

Improved Measurement of Neutrino Oscillations in the NOvA Experiment with a  
Decade of Data

Jozef Trokan-Tenorio

Milwaukee, Wisconsin

Master of Science, William and Mary, 2020  
Bachelor of Science, University of Wisconsin-Madison, 2018

A Dissertation presented to the Graduate Faculty of  
The College of William and Mary in Virginia in Candidacy for the Degree of  
Doctor of Philosophy

Department of Physics

The College of William and Mary in Virginia  
January 2026



## APPROVAL PAGE

This Dissertation is submitted in partial fulfillment of  
the requirements for the degree of

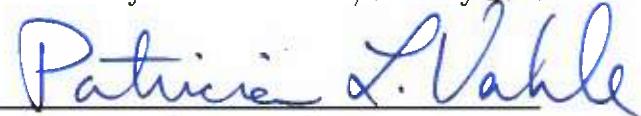
Doctor of Philosophy



---

Jozef Trokan-Tenorio

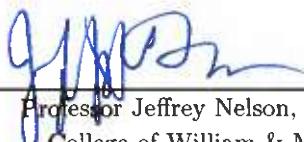
Reviewed by the Committee, January 2026



---

Committee Chair

Professor Patricia Vahle, Physics  
College of William & Mary



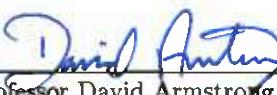
---

Professor Jeffrey Nelson, Physics  
College of William & Mary



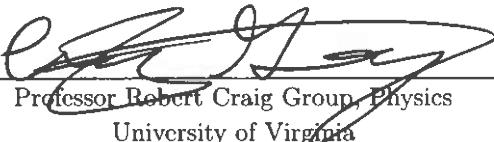
---

Professor Justin Stevens, Physics  
College of William & Mary



---

Professor David Armstrong, Physics  
College of William & Mary



---

Professor Robert Craig Group, Physics  
University of Virginia

## ABSTRACT

NOvA is a long-baseline neutrino oscillation experiment consisting of two functionally-identical tracking calorimeters, and a beam of neutrinos. The near detector is located at Fermilab, where it measures neutrinos coming from the NuMI beam. The beam can be run in neutrino or antineutrino mode, to produce a highly pure flux of muon (anti)neutrinos. The neutrinos then travel 810 km north to the much larger far detector, where we measure them again after they have oscillated. By measuring the appearance of electron (anti)neutrinos and the disappearance of muon (anti)neutrinos, we can make precise measurements of PMNS mixing matrix parameters, as well as the neutrino mass splitting  $\Delta m_{32}^2$ , and shed light on the remaining unknowns of neutrino mass ordering (the sign of  $\Delta m_{32}^2$ ),  $\delta_{\text{CP}}$ , and the octant of  $\theta_{23}$ . This dissertation presents a joint analysis of  $\nu_\mu \rightarrow \nu_\mu$  and  $\nu_\mu \rightarrow \nu_e$  data consisting of  $26.61 \times 10^{20}$  protons-on-target (POT) in neutrino mode and  $12.5 \times 10^{20}$  POT in antineutrino mode. Analysis improvements include new optimizations for our simulated light model, improved systematic uncertainties, a new sample of low-energy  $\nu_e$  events, and an additional focus on constraints from reactor neutrino experiments, including our first implementation of a constraint on  $\Delta m_{32}^2$ . Additional studies aimed at improving our sensitivity to the oscillation parameters are explored, including efforts to reclaim neutrino events that fail selections, the implementation of the reactor neutrino constraints, and significant improvements to Michel electron reconstruction. The best fit to the data falls in the normal mass ordering, and upper octant of  $\theta_{23}$ , with  $\Delta m_{32}^2 = 2.433_{-0.036}^{+0.035}$  ( $10^{-3}$  eV $^2$ ), and  $\sin^2 \theta_{23} = 0.546_{-0.075}^{+0.032}$ . Fitting with a constraint based on Daya Bay's measurement of  $\sin^2 2\theta_{13}$  gives us a  $1.36\sigma$  preference for the normal mass ordering, which increases to  $1.57\sigma$  when applying an additional constraint using their measurement of  $\Delta m_{32}^2$ . Our observed data falls in a region preferring combinations of the oscillation parameters that lead to similar oscillation probabilities for  $\nu_e$  and  $\bar{\nu}_e$ . Consequently we disfavor combinations of oscillation parameters that result in a large asymmetry, excluding  $\delta_{\text{CP}} = 0, \pi, 2\pi$  in the inverted mass ordering at  $> 3\sigma$ .

## TABLE OF CONTENTS

Acknowledgments	vi
List of Tables	vii
List of Figures	ix
Preface	1
Chapter 1. Theory of Neutrinos	3
1.1 Introduction	3
1.1.1 Particles in the Standard Model	5
1.1.2 Neutrino Interactions	6
1.2 Neutrino Oscillations	8
1.2.1 2-flavor oscillations	12
1.2.2 Matter Effects	14
1.2.3 3-flavor oscillations	15
1.3 Measuring Neutrino Oscillations	18
Chapter 2. The NOvA Experiment	22
2.1 Introduction	22
2.2 The NuMI Beam	23
2.2.1 The Proton Beam	24
2.2.2 The NuMI Target	24
2.2.3 The Focusing Horns	26

2.2.4	The Decay Pipe and Absorber	28
2.2.5	Off-axis detector location	29
2.2.6	Final Beam Components	30
2.3	The NOvA Detectors	34
2.3.1	From Cells to Full Detector	35
2.3.2	The Data Acquisition System	37
2.3.3	File organization	39
2.3.4	Event Display and Detector Differences	40
2.3.5	Data Quality Monitoring	43
2.3.6	Total Data Collected	45
Chapter 3. Detector Simulation & Calibration		46
3.1	Simulation	46
3.1.1	Beam Simulation	47
3.1.2	Detector Simulation	47
3.1.3	Cosmic and Rock Events	53
3.2	Calibration	53
3.2.1	Relative Calibration	55
3.2.2	Absolute Calibration	60
3.2.3	Accounting For Drift	61
Chapter 4. NOvA Software Infrastructure		65
4.1	Event Reconstruction	66
4.1.1	Hits	67
4.1.2	Clustering	68
4.1.3	$\nu_e$ reconstruction	69
4.1.4	$\nu_\mu$ Reconstruction	72
4.1.5	Cosmic Tracks	73

4.1.6	Michel electrons	74
4.2	Particle Identification	81
4.2.1	Event Classification with CVN	81
4.2.2	Muon Identification using ReMID	82
4.2.3	Prong CVN	83
4.2.4	Cosmic Rejection BDTs	84
4.2.5	Low-Energy $\nu_e$ BDT	84
4.3	Energy Estimation	85
4.3.1	$\nu_e$ energy estimation	85
4.3.2	$\nu_\mu$ energy estimation	86
Chapter 5.	3-Flavor Analysis Setup	89
5.1	Introduction	89
5.2	Event Selection	90
5.2.1	Spill and Data Quality	90
5.2.2	$\nu_e$ selection	91
5.2.3	$\nu_\mu$ selection	96
5.3	Analysis Binning	98
5.3.1	$\nu_e$ Binning	99
5.3.2	$\nu_\mu$ Binning	99
5.3.3	Transverse Momentum - $p_T$	101
5.4	Constructing Far Detector Predictions	102
5.4.1	Beam Decomposition	103
5.4.2	Extrapolation	109
5.4.3	Cosmic and Rock Backgrounds	112
5.4.4	Final FD Predictions	113
5.5	Systematic Uncertainties	116

5.5.1	Flux Systematics	116
5.5.2	Cross-Section Systematics	117
5.5.3	Detector Response and Calibration Systematics	119
5.5.4	Other Systematics	124
5.5.4.1	Normalization	125
5.5.4.2	Michel Tagging	125
5.5.5	Summary of uncertainties and the effect of extrapolation	126
5.6	Fitting Oscillation Parameters	128
5.6.1	Frequentist Techniques	129
5.6.2	Bayesian Techniques	131
Chapter 6.	Gaining Sensitivity	134
6.1	Daya Bay Constraints and Future Sensitivity	135
6.1.1	Implementing the constraints	135
6.1.2	Impact on oscillation parameters	137
6.1.3	Future Sensitivities	142
6.2	Reclaiming $\nu_\mu$ Events	148
6.3	Reclaiming $\nu_e$ Events	158
6.4	Michel electrons	167
6.4.1	Improving Simulation	167
6.4.2	Improving Reconstruction	169
6.4.3	Far Detector Michel electrons	175
6.4.4	Future improvements using Michel electrons	181
Chapter 7.	3-Flavor Analysis Results	188
7.1	The Observed Data	189
7.2	Fits and Results	195
7.2.1	Feldman-Cousins Corrections	205

7.2.2	The Effect of Reactor Neutrino Constraints	206
7.2.3	NOvA Results in a Global Context	211
Chapter 8.	Conclusions	215

## ACKNOWLEDGMENTS

First and foremost I would like to thank my advisor, Prof. Tricia Vahle for the constant guidance and insight into neutrino physics that has allowed me to develop my abilities as a researcher.

Thank you to the other members of my dissertation committee, Professors Jeff Nelson, Justin Stevens, David Armstrong, and Craig Group, for taking the time to read this dissertation.

Thank you to everyone on NOvA I have worked with. I have always felt welcomed and supported by the collaboration, in no small part because of the efforts of Young NOvA. I look forward to our new tradition of collaboration meeting karaoke, and more memorable nights out.

Thanks to all the other students and postdocs at William & Mary and at Fermilab, who have made this time enjoyable and fruitful, in and out of the office.

Thank you to all the educators in my life, past and present, who instilled in me a love of learning and a curiosity for the natural world. A particular thanks to Ms. Kiblawi, for proving that math is cool.

Thank you to my friends, for all the laughs and late nights on Discord that have kept me going through the most difficult times.

A special thank you to my family, my parents Mike and Rita, and my brother Chris for their endless love and unconditional support from day one.

Finally, a thank you to my partner Samantha, for sticking with me all these years, moving across the country with me (twice), and never wavering in your love and devotion.

## LIST OF TABLES

1.1	Summary of our current knowledge of the neutrino oscillation angles and mass splittings.	16
2.1	Summary of the predicted flavor profile of the NuMI beam in each running mode.	30
2.2	The primary decay modes of the mesons that produce the neutrinos for our beam, along with their corresponding decay fractions.	31
2.3	The decay modes of particles that contribute the most to the intrinsic electron neutrino backgrounds in our beam, along with their corresponding decay fractions.	31
3.1	Physics models in the GENIE version used for this analysis, along with the corrections applied to each model to improve data/MC agreement.	51
4.1	Energy biases and resolutions for $\nu_\mu$ events as well as individual $\mu$ and hadronic components.	88
5.1	Cosmic background events selected for all $\nu_e$ samples, and each $\nu_\mu$ quartile, before and after scaling to match beam livetimes.	113
5.2	Predicted number of additional events from rock interactions in the FD selected samples.	113
5.3	Predicted event counts for the $\nu_\mu$ , $\nu_e$ and Low-energy selected samples in the neutrino beam, and $\bar{\nu}_\mu$ and $\bar{\nu}_e$ samples in the antineutrino beam.	115

5.4	Name and short description of the effect in simulation for the 30 “large” cross section systematic knobs.	118
5.4	Name and short description of the effect in simulation for the 30 “large” cross section systematic knobs.	119
6.1	Tabulation of the largest sources of background in reconstructed Michel electrons at the ND.	174
6.2	Tabulation of the largest sources of background in reconstructed Michel electrons at the FD.	176
6.3	Michel electron tagging efficiency before and after updates to improve the simulation and reconstruction of Michels.	177
6.4	The number of true signal and wrong-sign $\nu_\mu$ events that contain a visible true Michel electron compared to all events.	182
6.5	The number of true signal and wrong-sign $\nu_e$ events that contain a visible true Michel electron compared to all events.	182
7.1	Predicted event counts for the $\nu_\mu$ , $\nu_e$ and Low-energy selected samples in the neutrino beam, and $\bar{\nu}_\mu$ and $\bar{\nu}_e$ samples in the antineutrino beam.	192

## LIST OF FIGURES

1.1 A summary of the classification of elementary particles in the Standard Model.	6
1.2 Feynman diagrams for charged (left) and neutral (right) current neutrino interactions.	7
1.3 Feynman diagrams for the four primary CC interaction processes we see in NOvA.	9
1.4 Neutrino and antineutrino CC cross sections divided by neutrino energy, as a function of neutrino energy, for different interaction types.	10
1.5 The oscillation probability in the two-flavor approximation, plotted as a function of baseline L with fixed energy, or as a function of energy E with a fixed baseline.	13
1.6 A schematic showing the two neutrino mass ordering possibilities, where the third neutrino mass state is either the heaviest (normal ordering) or the lightest (inverted ordering).	19
1.7 <b>Left:</b> Toy model of NOvA's observed far detector $\nu_\mu$ energy spectrum in the case where there would be no oscillations, vs. two examples with oscillations using different oscillation parameters. <b>Right:</b> The ratio of each of the oscillated spectra to the unoscillated spectra.	20
1.8 The probability of $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ vs. $\nu_\mu \rightarrow \nu_e$ oscillations for different combinations of the neutrino mass hierarchy, and octant of $\theta_{23}$ , with the value of $\delta_{\text{CP}}$ varying to trace out the ellipses.	21

2.1 <b>Left:</b> Side view depiction of the NOvA experimental setup (not to scale). <b>Right:</b> Top-down view showing the geographic location of the detectors as well as the off-axis location of the FD.	23
2.2 All of the main components of the NuMI beamline (not to scale).	23
2.3 Illustration of the main components of the Fermilab accelerator complex which provides protons for the NuMI beam.	25
2.4 Photograph of the NuMI target in its carrier.	26
2.5 Illustration showing potential paths taken by hadrons produced in the target and subsequently focused by Horn 1 and/or Horn 2.	27
2.6 Neutrino energy (a) and event rates (b) of neutrinos at the FD for different off-axis angles, assuming a baseline of 810 km.	30
2.7 FHC energy distributions for different neutrino types at the ND, broken down by parent meson.	32
2.8 RHC energy distributions for different neutrino flavors at the ND, broken down by parent meson.	33
2.9 Diagram showing the scale of the detectors, along with a zoom-in on the essential components.	34
2.10 Diagram of essential cell components.	36
2.11 Schematic showing what the X and Y views of the NOvA detector look like for an example neutrino event.	37
2.12 Overview of the Data Acquisition System for the FD.	38
2.13 A far detector event display, showing the high rate of cosmic events, as well as the effectiveness of isolating the beam-spill window.	41
2.14 A near detector event display, showing multiple neutrino interactions as well as rock events occurring in a single beam spill.	42

2.15 A plot of the Good Runs monitor, which indicates whether recent data collected is labeled as “Good” or is undergoing some type of failure mode.	44
2.16 Online Monitoring plots showing the FEB hit rate over a 24 hour period at the FD and ND.	44
2.17 Accumulated beam exposure at the FD over time.	45
3.1 Flowchart of our simulation chain.	46
3.2 Neutrino flux with PPFX corrections applied at the FD (top) and ND (bottom) for FHC and RHC beam modes.	48
3.3 Muon neutrino cross section on carbon for different interaction processes, as a function of energy.	49
3.4 Data/MC comparison of the reconstructed visible hadronic energy of $\nu_\mu$ CC selected events in the ND, for FHC and RHC beam modes.	50
3.5 A simple flowchart of the calibration procedure.	54
3.6 Bethe Bloch curve for muons in different materials.	55
3.7 Simulated PE as a function of W (position along the cell) for FD cells in the X view, split by different fiber brightness bins.	56
3.8 Block diagram illustrating the main steps in the relative calibration.	57
3.9 Example of the tricell hit condition used in calibration for a horizontal (Y-view) cell.	58
3.10 The results of the attenuation fit for a typical cell in the ND and FD.	60
3.11 FD plots showing the effect of relative calibration, averaged across all cells, with PE/cm for each view on top and PECorr/cm for each view on the bottom.	62
3.12 Simulated muon dE/dx, with the selected track window used for calibration shown.	63

3.13 The mean reconstructed/true energy shown for simulated FD hits in the Y-view before/after calibration is applied.	63
3.14 Observed PE over time for the FD, averaged over all X-view cells.	64
4.1 The typical topologies for $\nu_\mu$ CC (top), and $\nu_e$ CC (middle) events in our detectors, and an example NC (bottom) topology that appears $\nu_e$ -like.	66
4.2 Block diagram showing the order that reconstruction chain modules are run.	67
4.3 Illustration of the centroid-finding algorithm utilized by TDSlicer.	68
4.4 Flowchart of the $\nu_e$ reconstruction chain applied to an NC event.	70
4.5 A zoomed-in FD event display showing the Kalman tracks reconstructed for a simulated $\nu_\mu$ event.	73
4.6 Block diagram showing the sequence of steps in the MEFinder algorithm.	75
4.7 Normalized 2D template histograms of the Michel variables used in the MID calculation.	77
4.8 Data/MC comparison of the MID values of reconstructed SlcME and TrkME in the ND.	79
4.9 Data/MC comparison of calorimetric energy for reconstructed SlcME and TrkME in the ND for FHC events.	80
4.10 Pixel map of the Y-view of $\nu_e$ CC, and $\nu_\mu$ CC events along with the feature maps extracted for each.	82
4.11 Distribution of CVN output scores for $\nu_e$ CC and $\nu_\mu$ CC for simulated FD events, split by true interaction type.	83
4.12 Linear fits to the true energy (shown as a red line) for FHC and RHC events for the muon and hadronic components of the event.	87
5.1 Cut flow diagram for the $\nu_e$ selection at the FD.	93

5.2	The $\nu_e$ core sample cut optimization in 3D CoreBDT:CVNcosmic:CVNe space shown as CVNcosmic:CVNe projection for FHC and RHC with the optimal CoreBDT cut of 0.45 applied to both.	95
5.3	Cut flow diagram for the $\nu_\mu$ selection.	97
5.4	Predicted FD $\nu_e$ candidate events for the Low-energy, FHC and RHC samples.	100
5.5	Hadronic energy fraction vs. reconstructed $\nu_\mu$ energy for simulated events passing the full $\nu_\mu$ selection in the FD, for FHC and RHC.	101
5.6	A cartoon showing the different acceptances of the two detectors.	102
5.7	Reconstructed transverse momentum distributions for $\nu_\mu$ and $\nu_e$ events in the ND and FD.	102
5.8	Reconstructed $\nu_\mu$ energy distributions for the FHC contained and un-contained $\nu_\mu$ selected samples in the ND, used to constrain the pion and kaon flux in the BEN decomposition.	105
5.9	Simulated transverse vs. forward momentum for parent pions of $\nu_\mu$ and $\nu_e$ CC events in the ND, for FHC mode.	106
5.10	The number of reconstructed Michels per event in the ND $\nu_e$ selected sample, before corrections are applied, for all $\nu_e$ energies.	107
5.11	The ND $\nu_e$ selected samples in FHC and RHC after their corresponding decomposition methods have been applied.	108
5.12	The ND low-energy $\nu_e$ selected sample, before and after applying proportional decomposition.	108
5.13	Schematic diagram showing the extrapolation procedure for the $\nu_\mu \rightarrow \nu_\mu$ signal events.	109
5.14	Predictions for the FD $\nu_\mu$ selected events in the FHC and RHC samples, split by hadronic energy fraction, with extrapolation applied.	114

5.15 Predictions for the FD $\nu_\mu$ selected events in the FHC and RHC samples, with all hadronic energy bins summed, and with extrapolation applied.	114
5.16 Predicted FD $\nu_e$ candidate events for the FHC and RHC samples, with extrapolation applied.	115
5.17 Profile of the Cherenkov scaling parameter in the light level tune fit.	121
5.18 The ND and FD profiles of the X and Y $F_{\text{view}}$ scaling parameters in the light level tune fit.	121
5.19 Neutron prong energy for standard and MENATE-supplemented Geant4 simulation.	125
5.20 Selected cosmic events and their corresponding KDE for the FD FHC $\nu_\mu$ quartiles, before scaling to match beam livetimes.	126
5.21 Summary of errors on the integrated number of $\nu_\mu$ selected events for the FHC and RHC samples, with all quartiles summed.	127
5.22 Summary of errors on the integrated number of $\nu_e$ selected events for the FHC and RHC samples.	127
5.23 An illustration of the MCMC sampling technique used by NOvA, for a 1D example scenario.	133
6.1 Fake data fits showing frequentist confidence-level contours for $\Delta m_{32}^2$ vs. $\sin^2 \theta_{23}$ in the normal (top) and inverted (bottom) mass orderings, using several different constraints on the value of $\sin^2 2\theta_{13}$ .	138
6.2 Fake data fits showing frequentist confidence-level contours for $\sin^2 \theta_{23}$ vs. $\delta_{\text{CP}}$ in the normal (top) and inverted (bottom) mass orderings, using several different constraints on the value of $\sin^2 2\theta_{13}$ .	139

6.3	Fake data fits showing frequentist confidence-level contours for $\sin^2 \theta_{23}$ vs. $\delta_{\text{CP}}$ in the normal (top) and inverted (bottom) mass orderings, using nominal and $\Delta m_{ee}^2$ constraints.	140
6.4	Fake data fits showing frequentist confidence-level contours for $\Delta m_{32}^2$ vs. $\sin^2 \theta_{23}$ in the normal (top) and inverted (bottom) mass orderings, using nominal and $\Delta m_{ee}^2$ constraints.	141
6.5	The significance at which we could reject the incorrect mass hierarchy (left) or measure CP violation (right) as a function of the true value of $\delta_{\text{CP}}$ , for a future $63 \times 10^{20}$ POT level, with data split evenly between FHC and RHC.	142
6.6	The significance at which we could reject maximal mixing (left) and determine the octant of $\theta_{23}$ (right) as a function of the true value of $\theta_{23}$ , for a future $63 \times 10^{20}$ POT level, with data split evenly between FHC and RHC.	143
6.7	NOvA's sensitivity to determine the neutrino mass ordering at increasing levels of POT, for the normal (top) and inverted (bottom) mass orderings, as a function of the true value of $\delta_{\text{CP}}$ .	144
6.8	NOvA's sensitivity to discover CP violation at increasing levels of POT, for the normal (top) and inverted (bottom) mass orderings, as a function of the true value of $\delta_{\text{CP}}$ .	145
6.9	NOvA's sensitivity to reject maximal mixing at increasing levels of POT, for the normal (top) and inverted (bottom) mass orderings, as a function of the true value of $\delta_{\text{CP}}$ .	145
6.10	NOvA's sensitivity to determine the octant of $\theta_{23}$ at increasing levels of POT, for the normal (top) and inverted (bottom) mass orderings, as a function of the true value of $\delta_{\text{CP}}$ .	146

6.11 NOvA's sensitivity to determine the mass ordering (left) and measure CP violation (right) for the normal (top) and inverted (bottom) mass orderings, as a function of the true value of $\delta_{\text{CP}}$ .	147
6.12 NOvA's sensitivity to reject maximal mixing (left) and measure the octant of $\theta_{23}$ (right) for the normal (top) and inverted (bottom) mass orderings, as a function of the true value of $\theta_{23}$ .	148
6.13 Selection efficiency as a function of true energy for $\nu_\mu$ CC signal events, as well as mis-identified NC and $\nu_e$ CC background, for NOvA's 2020 analysis.	149
6.14 <b>Top Left:</b> All events failing the 2020 $\nu_\mu$ PID cut, for the FHC sample. <b>Top right:</b> All signal events failing the 2020 numu PID cut, for the FHC sample. <b>Bottom:</b> All events passing the full 2020 FHC $\nu_\mu$ selection.	150
6.15 Plots of simulated events from the 2020 analysis FHC sample failing the $\nu_\mu$ 2020 PID cut.	151
6.16 The figure of merit for a cut on the max Y and Z positions of the track, for events failing the 2020 $\nu_\mu$ PID selection.	152
6.17 Reconstructed energy distributions for the failed $\nu_\mu$ event sample before and after cutting on maximum track Y and Z position.	153
6.18 Reconstructed energy distributions for $\nu_\mu$ events failing the PID cut, split into the 3 bins of CVNm score.	154
6.19 Number of reconstructed Michel electrons present in the event, for $\nu_\mu$ events failing the PID cut, split into 3 bins of CVNm score.	155
6.20 1D sensitivity to $\Delta m_{32}^2$ and $\sin^2 \theta_{23}$ for a nominal fake data fit, versus 2 fits that include an additional sample of reclaimed $\nu_\mu$ events.	156
6.21 $1\sigma$ , $2\sigma$ , and $3\sigma$ contours in $\Delta m_{32}^2$ vs. $\sin^2 \theta_{23}$ space for fake data fits with/without an additional sample of reclaimed $\nu_\mu$ events.	157

6.22 True FD $\nu_e$ ( $\bar{\nu}_e$ ) events failing both the 2020 $\nu_e$ core and peripheral selections, for the 2020 analysis FHC (left) and RHC (right) samples.	159
6.23 2D confidence-level contours for $\sin^2 \theta_{23}$ vs. $\delta_{\text{CP}}$ in the normal (top) and inverted (bottom) mass orderings, with (left) and without (right) the additional FHC and RHC reclaimed $\nu_e$ samples.	159
6.24 A Venn diagram showing the number of true $\nu_e$ signal events that failed a particular combination of the 2020 $\nu_e$ selection cuts.	160
6.25 $\nu_e$ signal and beam background energy distributions, for the 4 Venn diagram categories containing the most $\nu_e$ signal.	161
6.26 The efficiency with which we select true $\nu_e$ CC events in the FD as a function of energy, for FHC (left) and RHC (right) samples.	163
6.27 Predicted asymmetry of FD $\nu_e$ and $\bar{\nu}_e$ events at the far detector.	164
6.28 Total FHC low-energy $\nu_e$ appearance events predicted as a function of $\delta_{\text{CP}}$ .	165
6.29 Comparison of 1D posterior density distributions for $\delta_{\text{CP}}$ from fits with and without the low-energy sample included.	166
6.30 ND Data/MC discrepancy seen in the $\Delta T$ of reconstructed Michel electrons in older simulation files.	168
6.31 Cartoon diagrams of the retrigerring logic used in the old simulation, and the correct version from the real detector firmware.	169
6.32 Data/MC comparison of Michel electron $\Delta T$ with old and new retrigerring logic for Michels that retrigger and those that don't.	170
6.33 ID template histograms using the old simulation.	171
6.34 ND MID template histograms using the new simulation.	172
6.35 Reconstructed energy, $\Delta T$ , and MID distributions for SlcME and TrkME in the near detector.	173

6.36 Reconstructed energy, $\Delta T$ , and MID distributions for SlcME and TrkME in the far detector.	178
6.37 Reconstructed energy of true SlcME Michel electrons in the FD, split by those that retrigger or don't retrigger the parent slice.	179
6.38 FD MID template histograms made with all simulation and reconstruction updates applied.	180
6.39 Different wrong-sign scaling scenarios for simulated FD FHC events with 0, 1, or 2+ reconstructed Michel electrons, for events passing the full FD $\nu_\mu$ selection.	184
6.40 Different wrong-sign scaling scenarios for simulated FD RHC events with 0, 1, or 2+ reconstructed Michel electrons, for events passing the full FD $\bar{\nu}_\mu$ selection.	185
6.41 Different wrong-sign scaling scenarios for simulated FD FHC events with 0, 1, or 2+ reconstructed Michel electrons, for events passing the full FD Core $\nu_e$ selection.	186
6.42 Different wrong-sign scaling scenarios for simulated FD RHC events with 0, 1, or 2+ reconstructed Michel electrons, for events passing the full FD Core $\bar{\nu}_e$ selection.	187
7.1 Observed and predicted energy spectra, for $\nu_\mu$ (left) and $\bar{\nu}_\mu$ (right) CC selected events in the FD.	189
7.2 Observed and predicted energy spectra, for $\nu_\mu$ (left) and $\bar{\nu}_\mu$ (right) CC selected events in the FD, split into four hadronic energy fractions but with the transverse momentum quantiles combined.	189

7.3 <b>Left:</b> Observed and predicted energy spectra, for $\nu_\mu$ CC selected events in the FD, for oscillated and unoscillated predictions. <b>Right:</b> Ratio of the observed data and oscillated prediction to the unoscillated prediction.	190
7.4 Observed and predicted energy spectra, for $\nu_e$ (left) and $\bar{\nu}_e$ (right) CC selected events in the FD.	191
7.5 Comparison of the predicted number of FD $\nu_e$ and $\bar{\nu}_e$ appearance events as a function of $\delta_{\text{CP}}$ in each mass ordering.	193
7.6 Bi-event plot showing total predicted appearance events in RHC and FHC for different combinations of octant, hierarchy, and $\delta_{\text{CP}}$ .	194
7.7 Frequentist 1,2,3 $\sigma$ confidence level contours for $\Delta m_{32}^2$ vs. $\sin^2 \theta_{23}$ from stats-only fits to the full FHC and RHC datasets, using the 1D Daya Bay constraint on $\sin^2 2\theta_{13}$ .	196
7.8 1D stats-only frequentist fits to the full FHC and RHC datasets, using the 1D Daya Bay constraint on $\sin^2 2\theta_{13}$ .	197
7.9 Frequentist 1,2,3 $\sigma$ confidence level contours for $\Delta m_{32}^2$ vs. $\sin^2 \theta_{23}$ from fits to the full FHC and RHC datasets, profiling over all systematic uncertainties, and using the 1D Daya Bay constraint on $\sin^2 2\theta_{13}$ .	199
7.10 1D frequentist fits to the full FHC and RHC datasets, profiling over all systematic uncertainties, and using the 1D Daya Bay constraint on $\sin^2 2\theta_{13}$ .	200
7.11 Systematic and statistical uncertainties on the measured values of $\Delta m_{32}^2$ , $\sin^2 \theta_{23}$ , and $\delta_{\text{CP}}$ at the best-fit point.	201
7.12 Frequentist 1,2,3 $\sigma$ confidence level contours for $\Delta m_{32}^2$ vs. $\sin^2 \theta_{23}$ from stats-only fits to the FHC-only $\nu_\mu + \nu_e$ dataset, using the 1D Daya Bay constraint on $\sin^2 2\theta_{13}$ .	203

7.13 1,2,3 $\sigma$ confidence level contours for $\Delta m_{32}^2$ vs. $\sin^2 \theta_{23}$ from stats-only fits to the $\nu_\mu$ -only dataset, using the 1D Daya Bay constraint on $\sin^2 2\theta_{13}$ .	204
7.14 Feldman-Cousins corrected 68% and 90% confidence level contours for the $\Delta m_{32}^2$ vs. $\sin^2 \theta_{23}$ space, found from a frequentist fit to the data.	205
7.15 <b>Left:</b> The best-fit oscillation parameters and FC-corrected 1 $\sigma$ uncertainties from the frequentist fit to the data, using the 1D Daya Bay constraint. <b>Right:</b> The significance with which we prefer the normal mass ordering when using the 1D or 2D Daya Bay constraints.	206
7.16 1D stats-only fits of $\Delta m_{32}^2$ using no reactor constraint, the 1D Daya Bay constraint, or the 2D Daya Bay constraint.	207
7.17 Marginalized posterior probability densities for $\Delta m_{32}^2$ in each mass ordering, from the Bayesian fit with systematics included, and for three different reactor constraint options.	208
7.18 1D stats-only frequentist fits of $\sin^2 \theta_{23}$ using no reactor constraint, the 1D Daya Bay constraint, or the 2D Daya Bay constraint.	209
7.19 2D Posterior probability density and Bayesian credible regions for $\sin^2 \theta_{23}$ vs. $\sin^2 2\theta_{13}$ marginalized jointly over the mass orderings.	210
7.20 Comparison of the $\Delta m_{32}^2$ vs. $\sin^2 \theta_{23}$ FC-corrected 90% confidence level contour from this analysis to the one from NOvA's 2020 analysis.	211
7.21 Comparisons of the 90% confidence level contours for $\Delta m_{32}^2$ in the normal MO between our new result and recent results from other experiments.	212
7.22 Comparison of central values and 1 $\sigma$ intervals for $\Delta m_{32}^2$ in the normal MO from various accelerator, reactor, and atmospheric experiments, and two joint fits.	213
7.23 New NOvA results for $\sin^2 \theta_{23}$ vs. $\delta_{\text{CP}}$ in the Normal and Inverted MOs, compared with previous NOvA results.	214

# Preface

This thesis describes the measurement of  $\nu_\mu \rightarrow \nu_\mu$  and  $\nu_\mu \rightarrow \nu_e$  neutrino oscillations in the NOvA experiment. Here we give an overview of the contents of each chapter, noting my relevant contributions to the analysis.

In Chapter 1 we cover the theory of neutrinos that is relevant to this dissertation, primarily a description of neutrino oscillations and how NOvA is sensitive to the parameters that govern them.

In Chapter 2 we describe the design of the NOvA experiment in detail, going over the NuMI neutrino beam and both the near and far detectors.

Chapter 3 describes how we produce a detailed simulation of our beam and detector to compare with the real data. We then describe the process of calibrating both the simulated and real data to convert measured quantities to meaningful physical units. As a calibration expert for NOvA, I performed these calibration procedures on the data that is used for this analysis.

Chapter 4 covers the various traditional and machine-learning algorithms that we use to process our raw data and identify the neutrino interactions. This involves reconstructing the topologies that particles make in our detector, and then using them to estimate the energies and flavor of the neutrino interaction that produced them. I made significant improvements to our Michel electron reconstruction technique described in this section, but the improvements will be targeted to future analyses, and are discussed in Chapter 6.

Chapter 5 details the 3-flavor analysis methodology, including the construction of our

predicted far detector samples, the various systematic uncertainties we incorporate, and the different Bayesian and frequentist fitting techniques used. I performed validation of the new low-energy sample and its implementation into our fitting framework. I also made improvements to the code that is used to generate the official prediction files for the analysis and produced many of them myself.

Chapter 6 covers various studies I performed that were aimed at gaining sensitivity to the oscillation parameters. I detail the implementation and validation of external constraints on our analysis using measurements from reactor neutrino experiments. I performed two studies aimed at reclaiming  $\nu_\mu$  and  $\nu_e$  events that failed our selection cuts. The  $\nu_e$  study also helped motivate the new low-energy sample which is discussed. Finally, I cover the detector simulation and Michel electron reconstruction improvements that I made, and discuss how they might be used to improve the 3-flavor analysis.

Chapter 7 presents the results of the joint 3-flavor fit of  $\nu_\mu$  and  $\nu_e$  data. I was a part of the frequentist analysis team and so the majority of results presented here come from that analysis. My contributions include the final observed energy spectrum plots for our data, the bi-event plots, and the frequentist fits presented without Feldman-Cousins corrections.

Finally, Chapter 8 summarizes the results and conclusions, and discusses future research directions .

# Chapter 1

## Theory of Neutrinos

### 1.1 Introduction

The neutrino was first postulated by W. Pauli in 1930 as a way to explain the continuous energy spectra that had been observed in  $\beta$  decays by J. Chadwick [1]. At the time,  $\beta$  decay was thought to be a two-body decay, and so a continuous energy spectrum would violate the principle of conservation of energy. By emitting a chargeless fermion along with the electron in the decay, a continuous energy spectrum would be allowed. In 1934 this idea was incorporated by E. Fermi into his theory of  $\beta$  decay, where he also gave the neutrino its name (Italian for “little neutral one”) [2].

Initial calculations of the small neutrino cross section led Pauli and others to believe that the neutrino would never be detected. However, in 1946 B. Pontecorvo proposed the idea of using nuclear reactors (which provide a large flux of neutrinos) combined with large detector volumes to achieve the necessary event rates [3]. F. Reines and C. Cowan were the first to succeed in using this method, when in 1956 they discovered the electron antineutrino using a liquid scintillator detector placed underground near a nuclear reactor [4]. The experiment measured the signal of a positron followed by a delayed signal from neutron capture in the inverse beta decay process:

$$p + \bar{\nu}_e \rightarrow n + e^+. \quad (1.1)$$

For the discovery of the electron neutrino, Reines was awarded the Nobel prize in 1995. Six years after the discovery of the  $\nu_e$ , the  $\nu_\mu$  was discovered at Brookhaven National Lab in 1962 by L. Lederman, M. Schwartz, and S. Steinberger [5], who pioneered the world's first accelerator-based neutrino beam. Notably, the design of the beam used to create neutrinos is the same general design used today. A high intensity beam of protons would hit a target, producing hadrons that then decay to neutrinos. The third generation lepton, the tau, was discovered in 1975 at the Stanford Linear Accelerator [6], and with it the Standard Model predicted a corresponding third generation of neutrino. However the discovery of the tau neutrino would not come until 2000. It was observed for the first time by the Direct Observation of the Nu Tau (DONUT) experiment at Fermilab [7], which used an emulsion cloud chamber to detect the short-lived tau signals.

While the three generations of neutrino were predicted by the Standard Model and subsequently discovered, the full picture was not yet complete, as the phenomena of neutrino oscillations presented a new puzzle. Pontecorvo had first proposed the idea of oscillations between neutrinos and antineutrinos in 1957 [8]. Z. Maki, M. Nakagawa, and S. Sagata expanded this idea further in 1962 to describe electron and muon neutrinos as combinations of mass eigenstates [9]. With this theory, the existence of neutrino oscillations would then also imply the neutrinos have mass. Pontecorvo then developed an intuitive theory of two-flavor neutrino oscillations in 1968 [10].

In that same year, R. Davis was measuring electron neutrinos from the sun using a tank of chlorine located underground in the Homestake mine in South Dakota. Neutrinos could capture on the chlorine, creating argon atoms, which he extracted to give a direct measurement of the solar neutrino flux [11]. However, he observed only a third of the events expected by the theory at the time. This became known as the solar neutrino problem. While some believed this indicated that the astrophysical estimates were wrong, many believed Davis had simply made an error in his measurements. However, other experiments also found deficits in the rate of solar neutrinos.

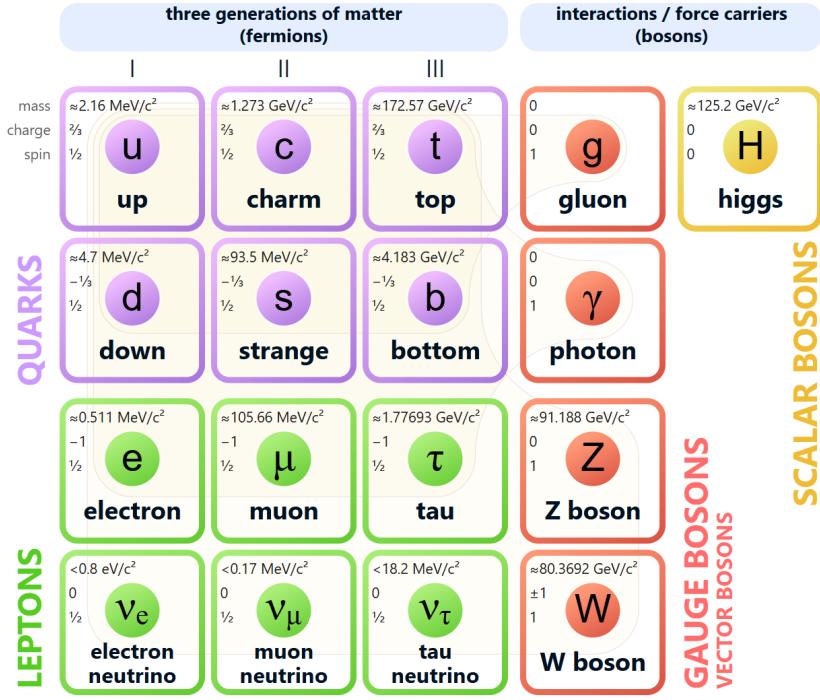
In 1998 the phenomena of oscillations was confirmed by the Super-Kamiokande (SK)

experiment [12]. They measured atmospheric neutrinos which result from the decays of pions, muons, and kaons produced in cosmic ray collisions in the atmosphere. Compared to a no-oscillation hypothesis, they found a deficit in the rates of muon neutrinos that traveled farther, through the Earth, on their way to the detector. Just a few years later the solar neutrino problem was also resolved, when in 2001 the Sudbury Neutrino Observatory (SNO) made a simultaneous measurement of the electron and total solar neutrino flux. The total solar flux agreed with the expected rate, while the  $\nu_e$  flux showed the deficit, confirming that the  $\nu_e$  were oscillating into other flavors [13]. For confirming that neutrinos oscillate and thus have mass, T. Kajita and A. McDonald, of the SK and SNO collaborations, respectively, received the Nobel prize in physics in 2015. Since then, the field has progressed to make precision measurements of the parameters describing neutrino oscillations.

### 1.1.1 Particles in the Standard Model

The Standard Model (SM) of particle physics is currently the most complete theory of fundamental particles and their interactions under the strong, electromagnetic, and weak forces. The classification of particles in the theory is summarized in Figure 1.1. There are two broad categories of particle, fermions that have  $\frac{1}{2}$  integer spin and obey the Pauli exclusion principle, and bosons that have integer spin. The gauge bosons give rise to each of the forces, with eight gluons mediating the strong force, the photon mediating the electromagnetic force, and the W and Z bosons mediating the weak force. The fermions are further classified based on their interactions and charges, into six quarks and six leptons (along with the antiparticles for each). Each of the six particles is paired up into three generations. The quarks carry both electric charge, color charge, and weak isospin, meaning they can interact via the electromagnetic, strong, and weak forces. Additionally, quarks are subject to confinement, meaning they only exist in color-neutral bound states of two or more quarks known as hadrons. Hadrons, containing an equal number of quarks and anti-quarks, are known as mesons, an example being pions (a first generation quark-antiquark

## Standard Model of Elementary Particles

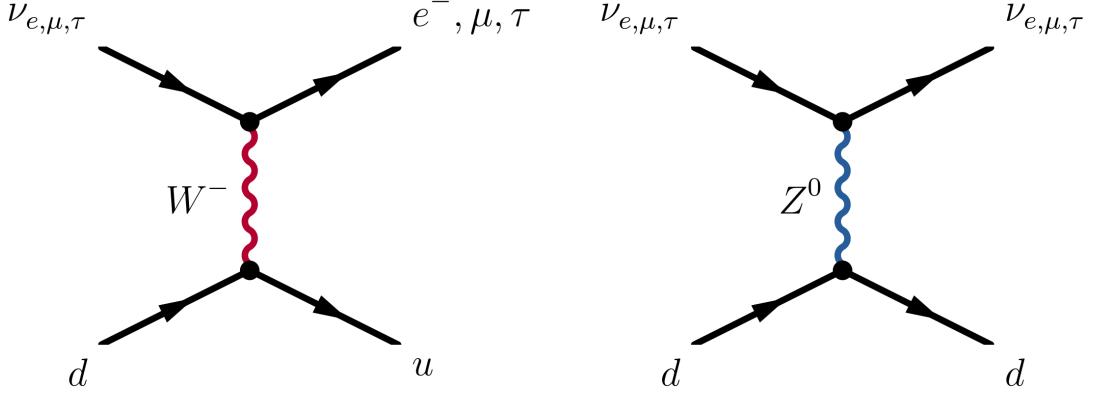


**Figure 1.1:** A summary of the classification of elementary particles in the Standard Model. From [14]

pair) or kaons (a first generation and strange quark combination). Three-quark states, such as the proton and neutron, are known as baryons. Each of the three generations of lepton are composed of a neutrino and charged lepton of the same flavor. The neutrinos do not have electric or color charge, and interact only via the weak force. The electron, muon, and tau all have an electric charge of -1, and interact via the weak and electromagnetic forces.

### 1.1.2 Neutrino Interactions

Within the Standard Model, neutrinos only interact via the weak force, undergoing either charged-current (CC) or neutral-current (NC) interactions. These are mediated by the  $W^\pm$  boson, and neutral  $Z$  boson, respectively. Figure 1.2 shows an example of Feynman



**Figure 1.2:** Feynman diagrams for charged (left) and neutral (right) current neutrino interactions. From [15].

diagrams for each of these. In a CC interaction, the neutrino exchanges a  $W^\pm$  boson, changing into its corresponding lepton pair ( $\nu_e$  to  $e$ ,  $\nu_\mu$  to  $\mu$ ,  $\nu_\tau$  to  $\tau$ ), which conserves lepton number. The sign of the  $W$  boson is chosen to conserve charge. In NC interactions the neutrino instead exchanges a neutral  $Z$  boson. Without the charge exchange, the neutrino remains in its initial state. In the context of inclusive  $\nu$  scattering, and neutrino oscillation experiments such as NOvA (the focus of this thesis, described in Chapter 2), the CC interactions we see can be generalized to the form

$$\overset{(-)}{\nu_l} + X \rightarrow l^{(+)} + X', \quad (1.2)$$

while the NC interactions can be generalized to

$$\overset{(-)}{\nu_l} + X \rightarrow \overset{(-)}{\nu_l} + X', \quad (1.3)$$

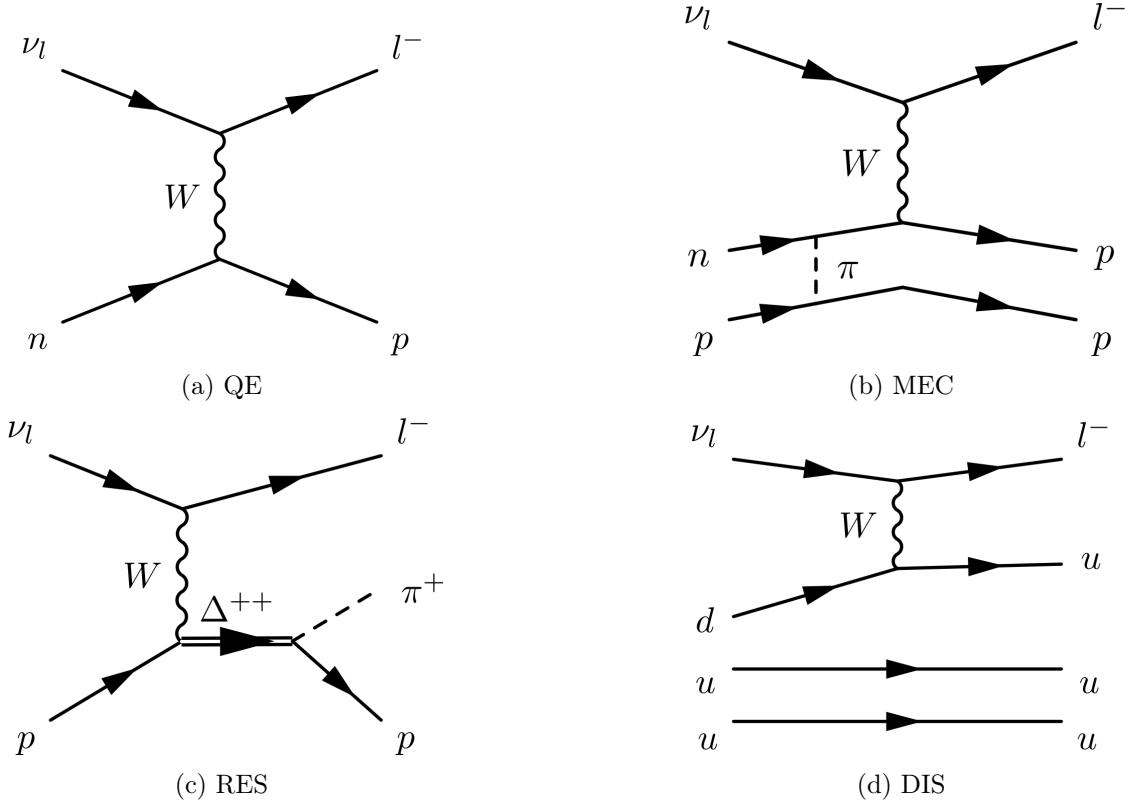
where  $X$  is the detector particle(s) that the neutrino interacted with,  $l$  is the charged lepton, and  $X'$  is the remaining final state particles [16]. Identifying the charged lepton in the final state of the CC interaction allows us to identify the flavor of the incoming neutrino. We cannot do such identification for NC interactions. The remaining final state particles  $X'$

in each interaction depend on the nature of  $X$ , the particle(s) that the neutrino interacted with, along with the energy of the incident neutrino. For example, the neutrino may interact with an entire nucleus, a single nucleon, or with enough momentum, individual constituent quarks. There are four main CC interaction processes that we see in NOvA: quasi-elastic (QE), meson-exchange current (MEC, also known as two-particle two-hole or 2p2h), resonant pion production (RES), and deep inelastic scattering (DIS). Figure 1.3 shows example Feynman diagrams for each of these processes. In the QE interaction (Fig 1.3a), the incoming neutrino interacts with a single nucleon. The charge exchange of the  $W$  boson leaves a different nucleon in the final state, and produces a charged lepton with the same flavor as the neutrino. In MEC interactions (Fig 1.3b), the neutrino interacts with a correlated pair of nucleons inside the nucleus. This leads to multiple nucleons in the final state as well. In RES interactions (Fig 1.3c) the exchanged boson has enough energy to create a  $\Delta$  resonance, which promptly decays to produce a pion in the final state in addition to a nucleon. Finally, DIS interactions (Fig 1.3d) involve a neutrino with enough energy to interact with an individual quark, breaking up the nucleon in the process and producing a variety of hadrons in the final state, known as a hadronic shower.

Each of these types of interaction has a different cross section that varies as a function of incident neutrino energy. The neutrino cross sections for QE, RES, and DIS interactions are plotted in Fig. 1.4. MEC events tend to occur between the energy ranges of QE and RES events [18]. NOvA primarily measures neutrinos in the 1-5 GeV range, so we see a mix of all kinds of interactions, primarily QE, MEC, and RES at lower energies, then RES and DIS at higher energies.

## 1.2 Neutrino Oscillations

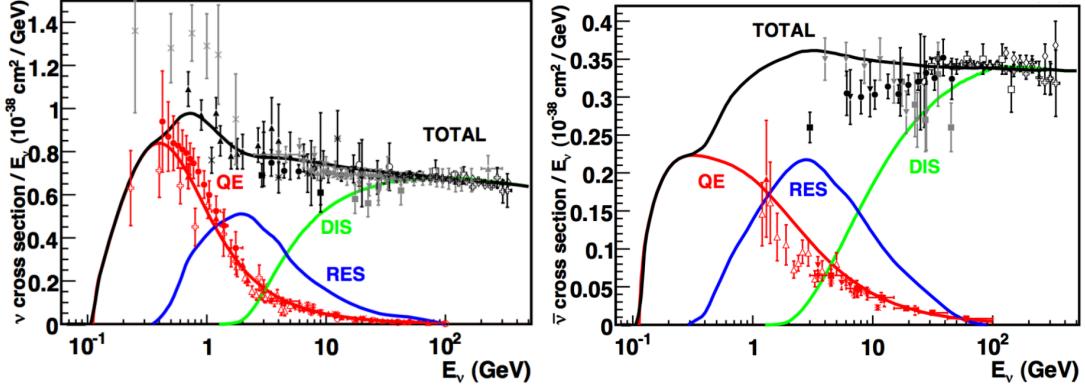
Neutrino oscillations are the phenomena whereby a neutrino created as one flavor can later be detected as a different flavor. This implies the neutrinos have separate mass and flavor eigenstates, which are superpositions of one another. This is in contrast to the Standard



**Figure 1.3:** Feynman diagrams for the four primary CC interaction processes we see in NOvA. From [17].

Model assumption that neutrinos are massless, and oscillations require modifying the theory. The flavor eigenstates represent the basis which diagonalizes the weak interaction, while mass eigenstates represent the basis that diagonalizes the free particle Hamiltonian. As the superposition propagates, the mass states will become out of phase with each other as their different masses cause them to propagate with different frequency. This causes a mixing of mass states that in turn results in a mixing of flavor states in the superposition. So, a neutrino created as a muon neutrino at a later time can be composed of a combination of electron, muon, and tau neutrino states. Since interactions occur in the flavor basis, it can then interact as one of the other flavors, and we would observe it as having oscillated.

Mathematically, we can write the neutrino flavor eigenstates ( $\nu_e$ ,  $\nu_\mu$ ,  $\nu_\tau$ ) as a linear



**Figure 1.4:** Neutrino (left) and antineutrino (right) per-nucleon CC cross sections (for an isoscalar target) divided by neutrino energy, as a function of neutrino energy. Solid lines show theoretical predictions for the total cross section as well as the different interaction types that contribute to the total, highlighted in different colors. Data points come from numerous experiments which are listed in the image source [19].

combination of mass eigenstates ( $\nu_1, \nu_2, \nu_3$ )

$$|\nu_\alpha\rangle = \sum_i U_{\alpha i}^* |\nu_i\rangle, \quad (1.4)$$

where  $|\nu_\alpha\rangle$  are the flavor eigenstates,  $|\nu_i\rangle$  are the mass eigenstates, and  $U_{\alpha i}^*$  is a  $3 \times 3$  unitary mixing matrix known as the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix [9, 10]. We can then derive the oscillation probabilities, using an approximation that the neutrinos are plane waves [20]. The full calculation requires a wave packet or quantum field theory treatment, however the plane wave approach will reach the same result for the purposes of this thesis.

Assume a neutrino of flavor  $\alpha$  is produced at time  $t_0$ . From Eq. 1.5 we can then write it as a superposition of the mass states that we have assumed to be plane waves with spatial momentum  $\mathbf{p}$  (and we are using natural units with  $\hbar = c = 1$ ),

$$|\nu_\alpha(t_0)\rangle = \sum_i U_{\alpha i}^* |\nu_i(\mathbf{p})\rangle. \quad (1.5)$$

The mass states are eigenstates of the free Hamiltonian:

$$\hat{H}|\nu_i(\mathbf{p})\rangle = E_i(\mathbf{p})|\nu_i(\mathbf{p})\rangle, \quad E_i(\mathbf{p})^2 = \mathbf{p}^2 + m_i^2. \quad (1.6)$$

The time evolution operator taking the state from  $t_0$  to  $t$  is given by  $e^{-i\hat{H}(t-t_0)}$ . The flavor state at time  $t$  is then

$$|\nu_\alpha(t)\rangle = e^{-i\hat{H}(t-t_0)}|\nu_\alpha(t_0)\rangle = \sum_i U_{\alpha i}^* e^{-iE_i(\mathbf{p})(t-t_0)}|\nu_i(\mathbf{p})\rangle. \quad (1.7)$$

We can then compute the probability that the state is in flavor  $\beta$  at time  $t$

$$P(\nu_\alpha \rightarrow \nu_\beta)(t) = |\langle \nu_\beta | \nu_\alpha(t) \rangle|^2 = \left| \sum_i U_{\beta i} U_{\alpha i}^* e^{-iE_i(\mathbf{p})(t-t_0)} \right|^2, \quad (1.8)$$

where we have applied the relation  $\langle \nu_i(\mathbf{p}) | \nu_j(\mathbf{p}) \rangle = \delta_{ij}$ .

The neutrinos in the NOvA experiment are ultra-relativistic which allows us to make the approximations

$$E_i(\mathbf{p})^2 - E_j(\mathbf{p})^2 \simeq \frac{1}{2} \frac{m_i^2 - m_j^2}{|\mathbf{p}|} + O(m^4), \quad L \simeq (t - t_0). \quad (1.9)$$

The final equation for the probability is then

$$P(\nu_\alpha \rightarrow \nu_\beta) = \sum_{i,j} U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^* e^{-i \frac{\Delta m_{ji}^2 L}{2|\mathbf{p}|}}, \quad (1.10)$$

where we have defined the mass splitting term  $\Delta m_{ij}^2 = m_i^2 - m_j^2$ . If we define  $W_{\alpha\beta}^{ij} = [U_{\alpha i} U_{\beta i}^* U_{\alpha j}^* U_{\beta j}]$  and apply the unitarity of the mixing matrix, we can rewrite the probability as

$$P(\nu_\alpha \rightarrow \nu_\beta) = \delta_{\alpha\beta} - 4 \sum_{j>i} \text{Re}[W_{\alpha\beta}^{ij}] \sin^2 \left( \frac{\Delta m_{ji}^2 L}{4E_\nu} \right) \pm 2 \sum_{j>i} \text{Im}[W_{\alpha\beta}^{ij}] \sin \left( \frac{\Delta m_{ji}^2 L}{2E_\nu} \right), \quad (1.11)$$

where the plus(minus) sign refers to neutrino(antineutrinos), and  $|\vec{p}| \sim E_\nu$ . In this form, we can clearly see how the probability “oscillates” sinusoidally as a function of  $L/E$ , hence the term oscillations. We can also see how the oscillations depend on the neutrinos having non-zero and non-degenerate masses. For the case where  $\beta = \alpha$  (the “survival probability”) this simplifies to

$$P(\nu_\alpha \rightarrow \nu_\alpha) = 1 - 4 \sum_{j>i} |U_{\alpha j}|^2 |U_{\alpha i}|^2 \sin^2 \left( \frac{\Delta m_{ji}^2 L}{4E_\nu} \right) \quad (1.12)$$

### 1.2.1 2-flavor oscillations

We can examine the oscillation probability in the simpler case of two-flavor oscillations. This is a useful approximation when considering how to design an oscillation experiment, as you will typically focus on a single oscillation channel ( $\nu_\mu \rightarrow \nu_e$ , for example). In this case, the PMNS matrix is a simple 2D rotation matrix

$$U_{\text{PMNS}} = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix} \quad (1.13)$$

and there is only a single mass splitting  $\Delta m^2$ . The oscillation probability becomes [20]

$$P(\nu_\alpha \rightarrow \nu_\beta) = \sin^2 2\theta \sin^2 \left( 1.27 \frac{\Delta m^2 (\text{eV}^2)}{E_\nu (\text{GeV})} L (\text{km}) \right), \quad \alpha \neq \beta \quad (1.14)$$

$$P(\nu_\alpha \rightarrow \nu_\alpha) = 1 - P(\nu_\alpha \rightarrow \nu_\beta), \quad (1.15)$$

where we have introduced factors of  $\hbar$  and  $c$  in order to convert to more familiar units. Equation 1.14 is known as the “appearance probability” since the observed final state is different than the initial, and Eq. 1.15 the “disappearance” or “survival” probability. The appearance probability is plotted in Fig. 1.5 as a function of the baseline  $L$  (distance between source and detector) and energy. We can see that it is a sinusoidal function of the

baseline, with a wavelength set by the oscillation length

$$L_{osc}(\text{km}) = 2\pi \frac{E_\nu(\text{GeV})}{1.27 \Delta m^2(\text{eV}^2)}. \quad (1.16)$$

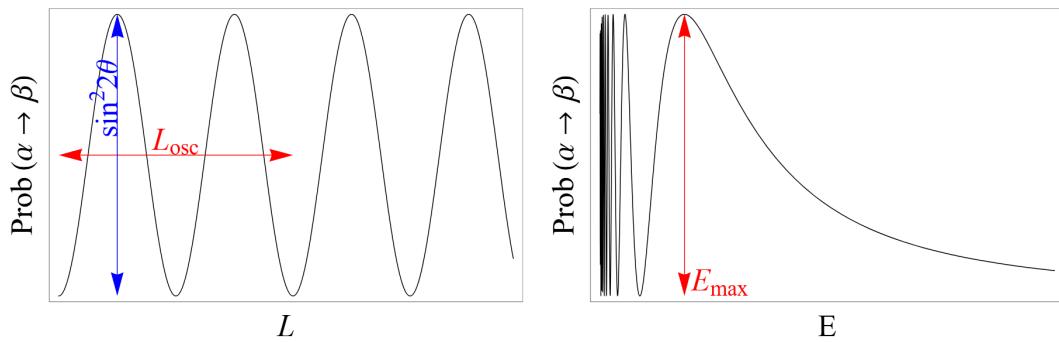
The amplitude of the oscillation is determined by the mixing angle in the leading  $\sin^2(2\theta)$  coefficient. It is maximal for  $\theta = \pi/4$ .

In experiments which have a fixed baseline, we instead typically measure the oscillation probability as a function of energy (RHS of Fig. 1.5). The peak of the first oscillation maximum (from right to left) is given by

$$E_{max}(\text{GeV}) = 1.27 \frac{\Delta m^2(\text{eV}^2) L(\text{km})}{\pi/2}, \quad (1.17)$$

which tells us about the mass splitting.

When designing a neutrino oscillation experiment in vacuum, the optimal settings will have the ratio of energy and baseline tuned to match the mass splitting you want to measure:  $E/L \sim \Delta m^2$ . In the case where  $E/L \ll \Delta m^2$  you are in the fast oscillation regime, where the experimental energy resolution is limiting. In the opposite case  $E/L \gg \Delta m^2$  the mass splitting and mixing angle cannot be disentangled [20].



**Figure 1.5:** The oscillation probability in the two-flavor approximation, plotted as a function of baseline  $L$  with fixed energy (left), or as a function of energy  $E$  with a fixed baseline (right). From [20].

### 1.2.2 Matter Effects

When neutrinos travel through matter, the oscillation probabilities are modified. The electron neutrinos can undergo coherent forward-scattering off the electrons present in the matter that is being traversed. This is known as the Mikheyev-Smirnov-Wolfenstein (MSW), or “matter”, effect and manifests as an additional potential applied to the  $\nu_e$  states [21, 22]

$$V_e = \pm \sqrt{2} G_F N_e. \quad (1.18)$$

Here  $G_F$  is Fermi’s constant,  $N_e$  is the electron number density, and the sign is positive for neutrinos, negative for antineutrinos. This results in an effective Hamiltonian

$$H_{\alpha\beta} = \sum_i U_{\beta i} \frac{m_i^2}{2E} U_{i\alpha}^* + A_{\alpha\beta} \quad (1.19)$$

where  $U$  is the PMNS matrix and the potential is introduced in the term  $A_{\alpha\beta} = \delta_{\alpha\beta} \delta_{\alpha e} V_e$ , which applies only to the electron neutrino states [15]. The solution to the Schrodinger equation with this effective Hamiltonian in the case of two-flavor oscillations can be written as Eq. 1.14 but with the  $\Delta m^2$  and  $\sin^2(2\theta)$  terms replaced by the effective terms [20]

$$\Delta m^2 \rightarrow \Delta m_M^2 = \sqrt{(\Delta m^2 \cos 2\theta - 2\sqrt{2} E G_F N_e)^2 + (\Delta m^2 \sin 2\theta)^2} \quad (1.20)$$

$$\sin^2 2\theta \rightarrow \sin^2 2\theta_M = \frac{(\Delta m^2 \sin 2\theta)^2}{(\Delta m_M^2)^2}. \quad (1.21)$$

In the limit where  $N_e$  goes to zero, this reduces to the vacuum oscillation probability.

From this we can see how the matter effects give us sensitivity to the sign of  $\Delta m^2$ . Note that the first term inside the square root in Eq. 1.20 depends on the relative signs of  $\Delta m^2$  and the potential term  $V_e$  which is positive (negative) for neutrinos(antineutrinos). Therefore the oscillation probability  $P(\nu_\mu \rightarrow \nu_e)$  is different for neutrinos versus antineutrinos, as well as for different signs of  $\Delta m^2$ .

### 1.2.3 3-flavor oscillations

We now consider the full 3-flavor oscillation probability in matter for NOvA. The  $3 \times 3$  PMNS matrix can be written in the form

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{CP}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{CP}} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & e^{i\alpha_{21}} & 0 \\ 0 & 0 & e^{i\alpha_{31}} \end{pmatrix} \quad (1.22)$$

$$= \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta_{CP}} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta_{CP}} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta_{CP}} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta_{CP}} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta_{CP}} & c_{23}c_{13} \end{pmatrix} \times \text{diag}(1, e^{i\alpha_{21}}, e^{i\alpha_{31}}) \quad (1.23)$$

where  $s_{ij} = \sin \theta_{ij}$  and  $c_{ij} = \cos \theta_{ij}$  with  $\theta_{ij}$  being the three mixing angles. Here  $\delta_{CP}$  is the CP-violating phase, and  $\alpha_{ij}$  are Majorana phases which do not enter into the oscillation probabilities and can be ignored. The oscillation probability (Eq. 1.11) is then determined by three mixing angles,  $\theta_{13}$ ,  $\theta_{12}$ ,  $\theta_{23}$ , two mass splittings  $\Delta m_{21}^2$  and  $\Delta m_{32}^2$ , and the CP violating phase  $\delta_{CP}$ <sup>1</sup>. The angles define the amplitude of the oscillation probability, and the mass splittings define the frequency of oscillation and position of the oscillation maximum as a function of L/E. Neutrino oscillation experiments will have different sources of neutrinos as well as different L/E that make them sensitive to different oscillation parameters. The main sources are solar neutrinos, neutrinos from nuclear reactors, neutrinos from particle accelerators (i.e. man-made neutrino beams), and atmospheric neutrinos from cosmic ray interactions in the atmosphere. To date, all of the mixing angles and mass splittings have been measured to a few percent uncertainty. Table 1.1 summarizes our current knowledge of the parameters and notes the type of experiments that are sensitive to each.

For NOvA, we are concerned with the disappearance oscillation channel  $P(\nu_\mu \rightarrow \nu_\mu)$  and the appearance oscillation channel  $P(\nu_\mu \rightarrow \nu_e)$ . We can simplify the subsequent

---

<sup>1</sup>It is also possible to define a third neutrino mass splitting  $\Delta m_{31}^2$ , which is related to the other two via the equation  $\Delta m_{32}^2 + \Delta m_{21}^2 = \Delta m_{31}^2$ . This may appear in some calculations.

Parameter	Approx. Value	Experiment Type
$\theta_{23}$	$45^\circ$	Atmospheric, Accelerator
$\theta_{13}$	$8.5^\circ$	Reactor, Accelerator
$\theta_{12}$	$33^\circ$	Solar, Reactor
$\Delta m_{32}^2$	$\pm 2.5 \times 10^{-3} \text{ eV}^2$	Accelerator, Reactor, Atmospheric
$\Delta m_{21}^2$	$7.5 \times 10^{-5} \text{ eV}^2$	Solar
$\delta_{\text{CP}}$	?	Accelerator

**Table 1.1:** Summary of our current knowledge of the oscillation parameters. Approximate values of the mixing angles and mass splittings derived from the Particle Data Group summary tables [23]. Each of the angles and mass splittings has been measured to a few percent uncertainty. The value of  $\Delta m_{32}^2$  has been measured but not the sign. The parameter  $\delta_{\text{CP}}$  has not yet been measured to a high degree of precision. The right column lists the types of neutrino oscillation experiments that are sensitive to a given parameter. For  $\theta_{13}$  the sensitivity primarily comes from reactor neutrino experiments, and for  $\theta_{12}$  the sensitivity primarily comes from solar experiments.

equations in this section by defining [16]

$$\Delta_{ij} \equiv \frac{\Delta m_{ij}^2 L}{4E}. \quad (1.24)$$

For  $P(\nu_\mu \rightarrow \nu_\mu)$  we start with Eq. 1.12 for the oscillation probability in vacuum

$$\begin{aligned} P(\nu_\mu \rightarrow \nu_\mu) &= 1 - 4 \sum_{j>i} U_{\mu j} U_{\mu j}^* U_{\mu i}^* U_{\mu i} \sin^2(\Delta_{ji}) \\ &= 1 - 4 \left( |U_{\mu 2}|^2 |U_{\mu 1}|^2 \sin^2(\Delta_{21}) + |U_{\mu 3}|^2 |U_{\mu 2}|^2 \sin^2(\Delta_{32}) + |U_{\mu 3}|^2 |U_{\mu 1}|^2 \sin^2(\Delta_{31}) \right) \\ &\approx 1 - 4 \left( c_{23}^2 s_{12}^2 c_{23}^2 c_{12}^2 \sin^2(\Delta_{21}) + c_{23}^2 c_{12}^2 s_{23}^2 s^2(\Delta_{32}) + c_{23}^2 s_{12}^2 s_{23}^2 s^2(\Delta_{31}) \right), \end{aligned} \quad (1.25)$$

where in the last step we have made an approximation based on current experimental knowledge that  $\theta_{13}$  is small relative to the other mixing angles [23], and so  $\sin^2 \theta_{13} \sim 0$  and  $\cos^2 \theta_{13} \sim 1$ . We can make a further approximation using the relationship between the mass splittings. We know  $|\Delta m_{21}^2|$  to be  $\sim 30$  times smaller than  $|\Delta m_{32}^2|$ , which lets

us approximate  $|\Delta m_{32}^2| \approx |\Delta m_{31}^2|$  [23]. We can then write Eq. 1.25 as

$$\begin{aligned} P(\nu_\mu \rightarrow \nu_\mu) &\approx 1 - 4 \cos^2 \theta_{23} \sin^2 \theta_{23} (\cos^2 \theta_{12} + \sin^2 \theta_{12}) \sin^2 (\Delta_{32}) \\ &= 1 - (2 \cos \theta_{23} \sin \theta_{23})^2 \sin^2 (\Delta_{32}) \\ &= 1 - \sin^2 2\theta_{23} \sin^2 (\Delta_{32}). \end{aligned} \quad (1.26)$$

To leading order we have recovered the result for the two-flavor approximation from Sec. 1.2.1. While the matter effects do play a role in the disappearance channel, it is known to be a very small correction compared to that for the  $\nu_e$  appearance channel [24]. Therefore for our purposes this approximation will suffice.

The full  $\nu_e$  appearance oscillation probability in vacuum, from Eq. 1.11, can be written as

$$P(\nu_\mu \rightarrow \nu_e) \approx P_{\text{atm}} + P_{\text{sol}} + 2\sqrt{P_{\text{atm}} P_{\text{sol}}} (\cos \Delta_{32} \cos \delta_{CP} \mp \sin \Delta_{32} \sin \delta_{CP}) , \quad (1.27)$$

where the  $\mp$  sign is minus (plus) for neutrinos (antineutrinos), and we have introduced the terms <sup>2</sup>

$$P_{\text{atm}} = \sin^2(\theta_{23}) \sin^2(2\theta_{13}) \sin^2(\Delta_{31}) \quad (1.28)$$

$$P_{\text{sol}} = \cos^2(\theta_{23}) \cos^2(\theta_{13}) \sin^2(2\theta_{12}) \sin^2(\Delta_{21}) . \quad (1.29)$$

As with the disappearance probability, we have approximated by dropping higher order terms in  $\sin^2 \theta_{13}$ . Including the matter effects, these equations become [25, 26]

$$P_{\text{atm}} = \sin^2(\theta_{23}) \sin^2(2\theta_{13}) \frac{\sin^2(\Delta_{31} \mp aL)}{(\Delta_{31} \mp aL)^2} \Delta_{31}^2 \quad (1.30)$$

$$P_{\text{sol}} = \cos^2(\theta_{23}) \cos^2(\theta_{13}) \sin^2(2\theta_{12}) \frac{\sin^2(\mp aL)}{\mp aL} \Delta_{21}^2 \quad (1.31)$$

---

<sup>2</sup>The subscripts *atm* and *sol* stand for "atmospheric" and "solar", after the  $\Delta_{31}$  and  $\Delta_{21}$  parameters that appear in each term which have been measured by atmospheric and solar neutrino experiments respectively.

where  $a = 2\sqrt{2}G_F N_e$  and the  $\mp$  signs are negative for neutrinos, and positive for antineutrinos. As described in Sec. 1.2.2, the relative signs of the matter potential terms  $a$  and the mass splittings will result in changes to the oscillation probability. From Eq. 1.27 we can also see how we get sensitivity to  $\delta_{\text{CP}}$  in the appearance channel. If  $\delta_{\text{CP}}$  is equal to 0 or  $\pi$ , then the last sine term is 0, and the overall oscillation probability in vacuum is the same for both neutrinos and antineutrinos, implying CP conservation. However, if  $\delta_{\text{CP}}$  takes on any other value, then CP is violated, and we get different oscillation probabilities for neutrinos and antineutrinos. This effect will compete with the matter effects when enhancing or suppressing the appearance rate for  $\nu_e$  ( $\bar{\nu}_e$ ) in NOvA, discussed further in the next section.

### 1.3 Measuring Neutrino Oscillations

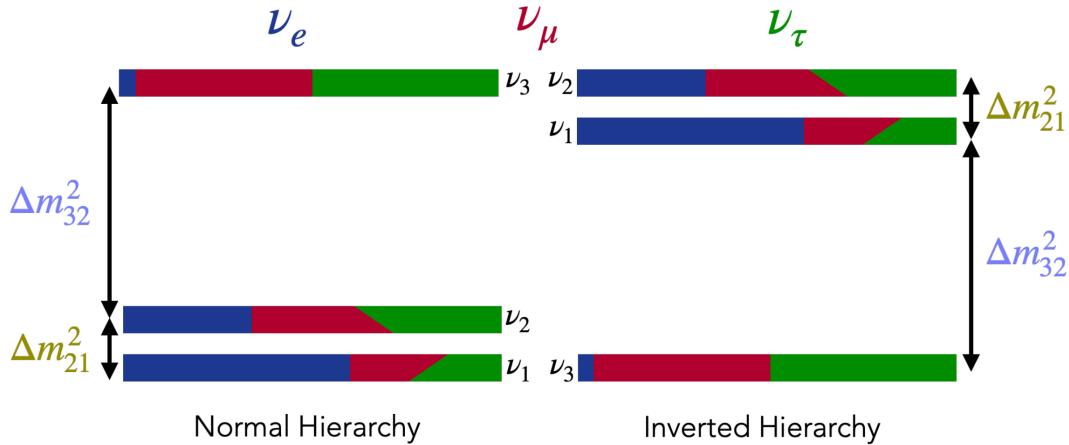
The goals of NOvA's 3-flavor oscillation analysis are to

- Measure the value and octant of  $\theta_{23}$
- Measure the magnitude of  $\Delta m_{32}^2$
- Constrain the values of  $\delta_{\text{CP}}$
- Determine the neutrino mass ordering (the sign of  $\Delta m_{32}^2$ ).

The octant of  $\theta_{23}$  refers to whether  $\theta_{23}$  is greater than  $\pi/4$  (upper octant) or less than  $\pi/4$  (lower octant).  $\theta_{23} = \pi/4$  is referred to as maximal mixing. If the mixing is maximal then it would imply the  $\nu_3$  mass state has an equal mix of  $\nu_\mu$  and  $\nu_\tau$  states, hinting at some fundamental symmetry in the lepton sector.

The neutrino mass ordering (also sometimes called the neutrino mass hierarchy) refers to the question of whether the third neutrino mass state is the heaviest or the lightest. In oscillation measurements we don't measure the neutrino mass directly, only the mass splittings. It is known that the solar mass splitting  $\Delta m_{21}^2$  is both small and positive,

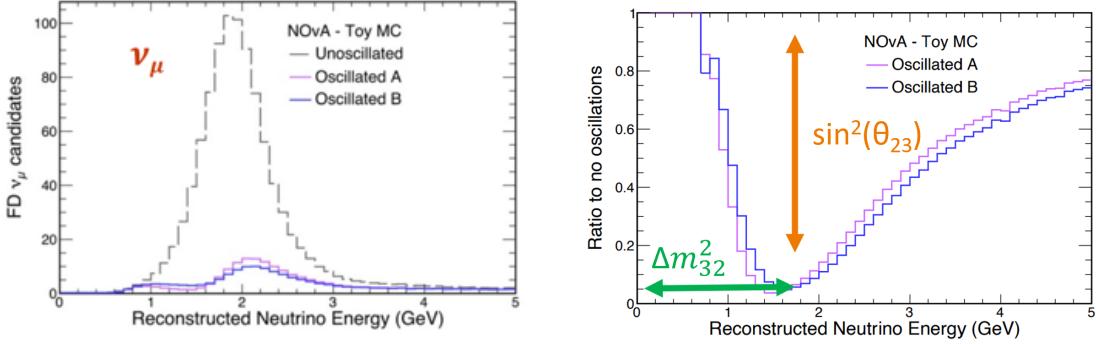
implying that  $\nu_1$  and  $\nu_2$  are of similar mass, and  $m_2$  is heavier than  $m_1$ . However, the sign of  $\Delta m_{32}^2$  is unknown. This leads to two possible orderings of the mass states, shown in Fig. 1.6. For the “normal” mass ordering, the  $\nu_3$  state is the heaviest, which results in a positive  $\Delta m_{32}^2$ . For the “inverted” ordering,  $\nu_3$  is the lightest, which results in a negative  $\Delta m_{32}^2$ .



**Figure 1.6:** A schematic showing the two neutrino mass ordering possibilities, where the third neutrino mass state is either the heaviest (normal ordering) or the lightest (inverted ordering). Colors show the rough flavor composition of each state, with the diagonals representing varying values of  $\delta_{\text{CP}}$ . From [16].

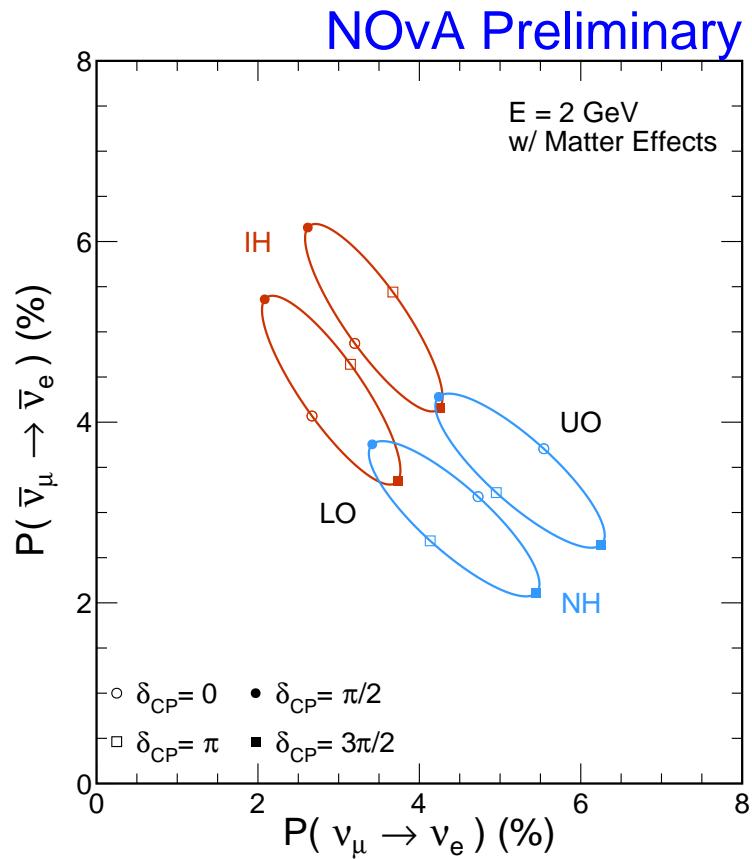
The two oscillation channels we measure probe the parameters in different ways. In the  $\nu_\mu$  disappearance channel  $P(\nu_\mu \rightarrow \nu_\mu)$  we observe the disappearance of  $\nu_\mu$  ( $\bar{\nu}_\mu$ ) as they travel from our beam source to the far detector and oscillate into other flavors. Fig. 1.7 shows a toy model of the far detector  $\nu_\mu$  energy spectra for two different values of oscillation parameters, compared to a no-oscillations case. Taking the ratio of the oscillated to unoscillated spectra, we get a characteristic dip shape. Measuring the location and depth of the minimum will tell us the values of  $\Delta m_{32}^2$  and  $\sin^2 \theta_{23}$  respectively. However, the oscillation probability is symmetric about  $\theta_{23} = \pi/4$ , which prevents us from determining the octant of  $\theta_{23}$  by this channel alone. To make that measurement, along with the other goals, we turn to the appearance channel.

In the  $\nu_e$  appearance channel  $P(\nu_\mu \rightarrow \nu_e)$ , we are sensitive to the value of  $\delta_{\text{CP}}$ , the



**Figure 1.7:** **Left:** Toy model of NOvA’s observed far detector  $\nu_\mu$  energy spectrum in the case where there would be no oscillations, vs. two examples with oscillations using different oscillation parameters. **Right:** The ratio of each of the oscillated spectra to the unoscillated spectra. The location and depth of the minimum are determined by the oscillation parameters. From [27].

octant of  $\theta_{23}$ , and the mass ordering. The octant of  $\theta_{23}$  will lead to an overall enhancement or suppression of  $\nu_e$  ( $\bar{\nu}_e$ ) appearance. As discussed in the previous section, the matter effects and CP violation lead to differences in the appearance probability for  $\nu_e$  ( $\bar{\nu}_e$ ) events, and this difference changes when considering the normal or inverted mass orderings. From the matter effects, in the normal mass ordering, we see an enhancement of  $\nu_e$  appearance and a suppression of  $\bar{\nu}_e$  appearance. In the inverted ordering, we see an enhancement of  $\bar{\nu}_e$  appearance and a suppression of  $\nu_e$ . How these combine with the effects from the value of  $\delta_{\text{CP}}$  can be visualized in Fig. 1.8. This shows the appearance probability for  $\bar{\nu}_e$  vs. the appearance probability for  $\nu_e$  at NOvA’s baseline and energy, for different combinations of the mass ordering, value of  $\delta_{\text{CP}}$ , and octant of  $\theta_{23}$ . We can see that for certain combinations of  $\delta_{\text{CP}}$  and mass ordering, the enhancements/suppressions work in the same direction, leading to a large asymmetry in the probability between  $\nu_e$  and  $\bar{\nu}_e$  appearance (normal mass ordering with  $\delta_{\text{CP}} = 3\pi/2$  for example). For other combinations, the effects work in opposite directions, canceling out to give a similar appearance probability for  $\nu_e$  and  $\bar{\nu}_e$ .



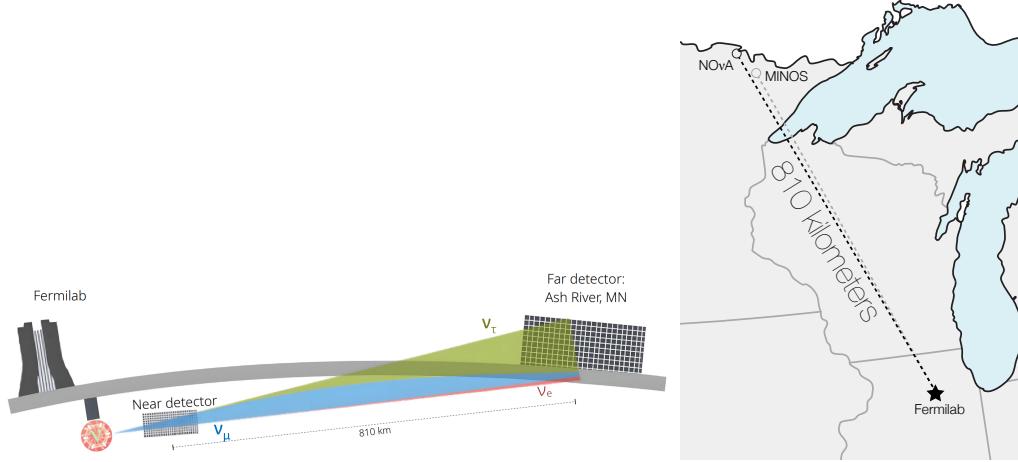
**Figure 1.8:** The probability of  $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$  vs.  $\nu_\mu \rightarrow \nu_e$  oscillations for different combinations of the neutrino mass hierarchy, and octant of  $\theta_{23}$ , with the value of  $\delta_{CP}$  varying to trace out the ellipses. The normal hierarchy (NH) is shown in blue and the inverted hierarchy (IH) in red, with  $\theta_{23}$  values in the lower and upper octants on the left (LO) and right (UO). From [28].

## Chapter 2

# The NOvA Experiment

### 2.1 Introduction

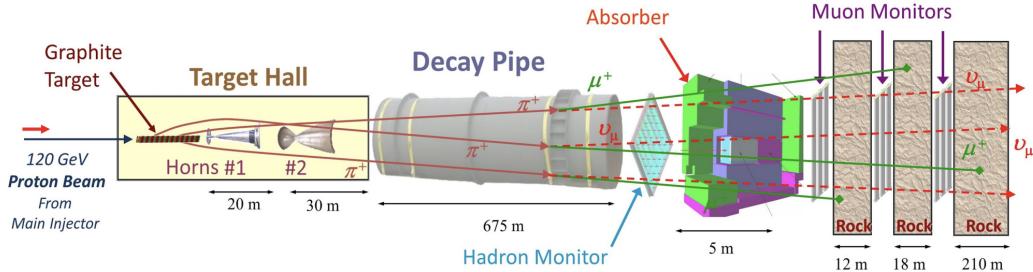
The NuMI Off-Axis Neutrino Experiment (NOvA) is a long-baseline neutrino oscillation experiment based at Fermilab in Batavia, IL. Figure 2.1 shows the experimental setup. It uses two functionally identical detectors separated by a baseline of 810 km to measure the oscillation of neutrinos produced by the Neutrinos at the Main Injector (NuMI) beam. The smaller Near Detector (ND) is located 105 m underground, 1 km from the beam source at Fermilab, and samples the neutrino interactions immediately after they are produced. The neutrinos then travel through the Earth 810 km north to the much larger Far Detector (FD) located on the surface in Ash River, Minnesota, where their interactions are sampled again after oscillating. The beam is primarily made up of muon neutrinos, and the oscillation channels that we measure are  $\nu_\mu$  ( $\bar{\nu}_\mu$ )  $\rightarrow$   $\nu_e$  ( $\bar{\nu}_e$ ) and  $\nu_\mu$  ( $\bar{\nu}_\mu$ )  $\rightarrow$   $\nu_\mu$  ( $\bar{\nu}_\mu$ ). Measuring the disappearance of muon neutrinos and the appearance of electron neutrinos lets us measure the mixing angle  $\theta_{23}$ , the neutrino mass-splitting  $\Delta m_{32}^2$ , and the CP-violating phase  $\delta_{\text{CP}}$ . Both detectors are located slightly off-axis of the beam to optimize the measurement of the mass ordering (the sign of  $\Delta m_{32}^2$ ). This chapter describes the experimental design of NOvA, starting with the beam and then the detector designs, including their off-axis placement.



**Figure 2.1:** Left: Side view depiction of the NOvA experimental setup (not to scale) [29]. Note the location of the near detector and beam source underground at Fermilab, the far detector aboveground, and the path the beam takes through the earth. Right: Top-down view showing the geographic location of the detectors as well as the off-axis location of the FD [30].

## 2.2 The NuMI Beam

Neutrinos at the Main Injector (NuMI) is used to create the neutrino beam source for NOvA. It works by impinging a 120 GeV beam of protons on a graphite target, producing hadrons (including charged pions and kaons), that are then focused before decaying to produce neutrinos, primarily through the  $\pi^{+(-)} \rightarrow \nu_\mu (\bar{\nu}_\mu) + \mu^{+(-)}$  channel. The beamline components are illustrated in Fig 2.2.



**Figure 2.2:** All of the main components of the NuMI beamline (not to scale). Particles shown correspond to the Forward Horn Current (FHC) configuration. Image from [31].

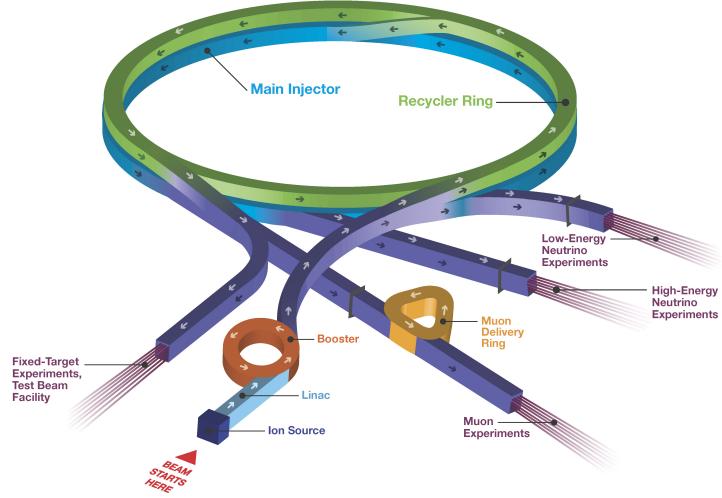
### 2.2.1 The Proton Beam

The Fermilab accelerator complex that produces the proton beam is illustrated in Fig. 2.3. The protons initially start out as  $H^-$  ions from an ion source. They are produced with an energy of 35 keV and are immediately accelerated to 400 MeV by a linear accelerator called the Linac. They then enter the first accelerator/storage ring called the Booster. Here, the electrons are stripped off using a carbon foil, leaving only protons, which are accelerated further to 8 GeV. A kicker magnet then sends the protons in 1.6  $\mu s$ -long batches to the larger storage ring called the Recycler Ring, which sits directly on top of the final accelerator, the Main Injector. The Main Injector and Recycler have a circumference  $7\times$  that of the Booster, so they can accommodate 6 batches at a time, with some overhead for the kicker's rise and fall time. The Recycler uses a process called “slip-stacking” which allows two batches to be combined for double the intensity [32]. After accelerating the 12 (6 pairs of slip-stacked) batches of protons to their full energy of 120 GeV, they are then sent to their final destination in the NuMI target hall. Each set of 12 batches is referred to as one “spill”, and delivers up to  $\sim 5 \times 10^{13}$  protons to the target over the course of 10  $\mu s$  [32]. During normal running conditions the time between spills is  $\sim 1$  second. The total number of protons delivered is referred to as the amount of protons-on-target (POT) and is used as the metric for how much data we've accumulated. The beam power varied significantly over the course of NOvA's data taking as upgrades to beamline components were implemented, increasing the amount of POT per spill, and reducing the time between spills. The beam ranged from  $\sim 500$  kW up to a recent record of 1 MW, with an average beam power of  $\sim 700$  kW.

### 2.2.2 The NuMI Target

The production target of the NuMI beam is a series of 48 thin graphite fins. Each fin is 25 mm long in the beam direction, and 9 mm across. The fins are spaced 0.3 mm apart, giving a net target length of  $\sim 120$  cm, or approximately two interaction lengths [34].

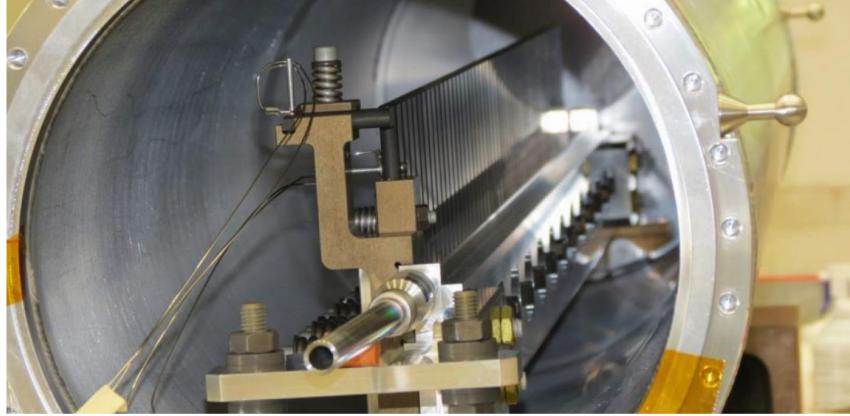
**Fermilab Accelerator Complex**



**Figure 2.3:** Illustration of the main components of the Fermilab accelerator complex which provides protons for the NuMI beam. Image from [33].

Figure 2.4 shows an image of the target, while the target’s placement in the full NuMI beamline can be seen in Fig. 2.2. The target is long in order to increase the number of hadrons produced, and thin to allow the produced hadrons to escape out the sides without undergoing too many secondary interactions in the target, thus maximizing the neutrino flux. The thinness of the target is balanced against the requirement that it be robust enough to endure the high intensity proton beam, which requires a large enough target so that the energy deposited per unit volume is not too high, with a maximum simulated temperature of 913° C [34]. To help reduce the heat generated by these depositions, the fins have been brazed onto two stainless steel pipes that run along the beam direction and provide liquid-based cooling [35].

Just upstream of the target sits a collimating baffle that protects the more sensitive downstream components from any misdirected proton beam. It is composed of a cylindrical graphite rod 57 mm in diameter and 1.5 m long, encased in a thin aluminum tube,

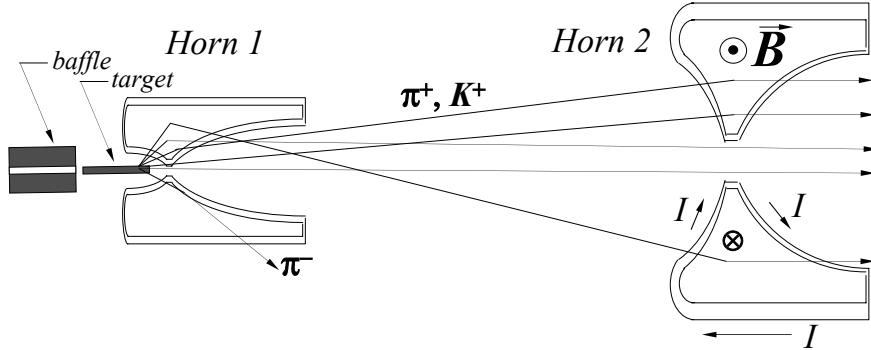


**Figure 2.4:** Photograph of the NuMI target in its carrier. The individual fins can be seen on top with one of the liquid cooling pipes running below. Image from [36].

with an 11 mm diameter hole running down the center for the beam to pass through. Thermocouples attached to the downstream end of the baffle can detect any misdirected beam as a sudden increase in temperature. It is able to withstand several spills of the full intensity proton beam, long enough for the misdirected beam to be detected and shut off. The beam is designed to be operated with up to 3% of it hitting the inner wall of the baffle, but under normal operation it sees less than 1% [35].

### 2.2.3 The Focusing Horns

Downstream of the target are two magnetic “horns” that act as lenses to focus hadrons produced at an angle to the beam axis, thus increasing the resulting neutrino flux in the forward direction. They are made of aluminum, with the inner conducting surface of revolution in the shape of a parabola. Current is run through the horns to produce a toroidal magnetic field, which focuses either positively or negatively charged hadrons depending on the direction of the current. Any hadrons that are under or over-focused in the first horn are refocused by the second horn, and particles of the opposite charge are ejected. Figure 2.5 illustrates this concept. We call the mode that focuses positively charged hadrons the Forward Horn Current (FHC) mode, while negatively charged hadrons are focused in Reverse Horn Current (RHC) mode. The primary decay mode is  $\pi^{+(-)} \rightarrow \nu_\mu + \mu^{+(-)}$ , so



**Figure 2.5:** Illustration showing potential paths taken by hadrons produced in the target and subsequently focused by Horn 1 and/or Horn 2. Hadrons that are over or under-focused by the first horn are re-focused onto the beam axis by the second horn. Particles produced directly on-axis are unaffected. This diagram shows the Forward Horn Current (FHC) configuration. For Reverse Horn Current the current would flow in the opposite direction, flipping the magnetic field direction, and instead focusing negative pions and kaons, while ejecting positive ones. Note that this diagram has the target inside horn 1 (the MINOS configuration) while for NOvA it is pulled back to be just outside horn 1, but the focusing concept is the same. Image from [37].

FHC corresponds to a neutrino beam, and RHC corresponds to an antineutrino beam. In either mode, if a particle is traveling directly along the beam axis, it will not be affected by the magnetic fields, and can pass through. This will lead to a small amount of  $\bar{\nu}_\mu$  ( $\nu_\mu$ ) being produced in the FHC(RHC) beam. We refer to these unwanted opposite-sign neutrinos as “wrong-sign” neutrinos, and they constitute one of our irreducible beam backgrounds.

Adjusting the position of the horns and target will change the resulting energy distribution of produced neutrinos. Lower energy hadrons will tend to contain a higher relative transverse momentum, and so will be traveling at higher angles relative to the beam direction. Thus, placing the first horn closer to the target will capture more of the low-energy hadrons, and in the case of the so-called “Low-Energy” configuration where the target is placed entirely inside Horn 1, all angles are captured. Moving either Horn further away from the target will then select for higher-energy hadrons and neutrinos. For NOvA, a “Medium-Energy” configuration was selected, where Horn 1 is placed 15 cm from the down-

stream end of the target, and Horn 2 is placed 10 m from the upstream end of Horn 1. The resulting energy distribution was shown to give the best sensitivity to the mass ordering when considering both FHC and RHC running [38].

#### 2.2.4 The Decay Pipe and Absorber

After the hadrons have been focused by the horns they travel through a large steel pipe (the decay pipe) to facilitate their decay and production of neutrinos. The pipe is 2 m in diameter and 675 m long, and is placed 50 m downstream of the target. Initially the inside was vacuum until corrosion was discovered on one of the beam windows. It was then filled with helium to slightly below atmospheric pressure. This relieved some pressure on the beam window to ensure it didn't fail, while preventing further corrosion and maintaining a low density environment for the particles to decay in. Just beyond the end of this pipe sits the hadron absorber, a large box-shaped structure 5.5 m wide  $\times$  5.6 m tall  $\times$  8.5 m long, with an aluminum core surrounded by layers of steel and concrete. It serves to absorb most of the remaining non-neutrino particles in the beam, preventing unwanted particles from reaching the near detector, as well as protecting groundwater and personnel from excess radiation [35]. The hadron absorber is instrumented upstream with a  $7 \times 7$  grid of helium-filled ionization chamber detectors, with a total area of 1 m  $\times$  1 m. Called the "hadron monitor", it measures the residual hadron flux, and is used to gauge the quality of the NuMI target and track alignment of the target and horns to within  $\pm 0.1$  mm [31]. Muons and, of course, neutrinos are able to pass through the absorber. There is 240 m of dolomite rock in-between the end of the NuMI beamline and the ND hall, with additional ionization chamber detectors placed in alcoves in the rock downstream of the absorber. These "muon monitors" are slightly larger with an area of 2.3 m  $\times$  2.3 m, and measure the two-dimensional profile of muons at different energies. As we move downstream, lower-energy muons are ranged out by the rock, and so the monitors have increasing muon detection thresholds of 4, 10, and 20 GeV [35]. Observing these muons provides an additional measure of the quality and intensity of the beam. The positions of

the absorber, monitors, and rock in the beamline can be seen in Fig. 2.2. The last 210 m of rock serves to range out the majority of remaining muons, and we are left with our final beam of neutrinos.

### 2.2.5 Off-axis detector location

Another important aspect of the NOvA experiment is the placement of the detectors 14.6 mrad ( $0.865^\circ$ ) away from the central beam axis. This is done to optimize the  $\nu_e$  appearance channel, given the two-body decay kinematics that produce the majority of our neutrinos. In their own rest frames, the pions and kaons decay isotropically, giving mono-energetic neutrinos. However, in the lab frame, the pions decay-in-flight, and the neutrinos are Lorentz-boosted. In the case of  $\pi \rightarrow \nu + \mu$ , the energy of a neutrino with a small angle  $\theta$  relative to the pion direction is given by

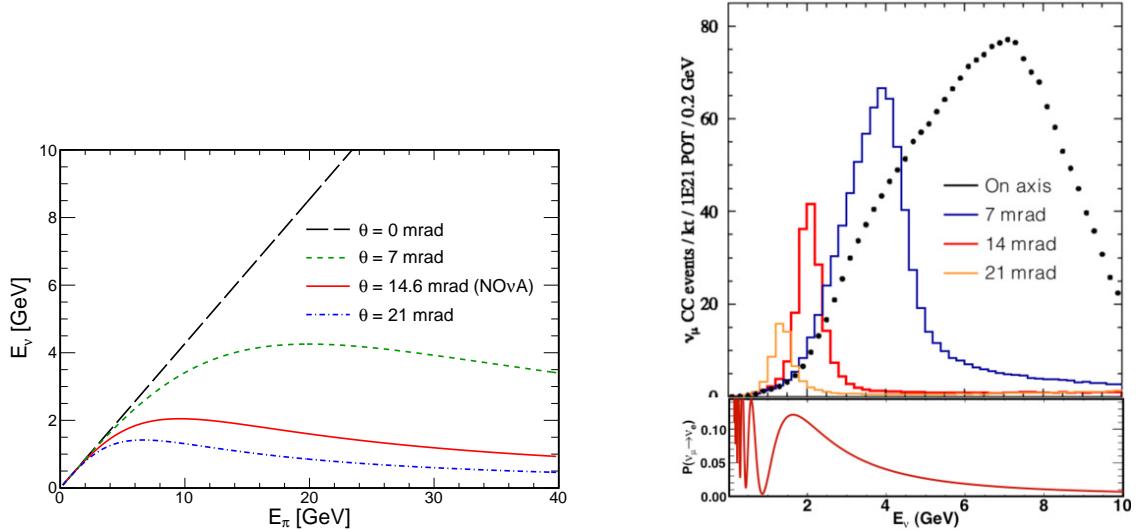
$$E_\nu = \frac{0.43E_\pi}{1 + \gamma^2\theta^2} \quad (2.1)$$

where  $E_\pi$  is the energy of the pion,  $\gamma$  is  $\frac{E_\pi}{m_\pi}$  and  $m_\pi$  is the mass of the pion [38]. For a detector with area  $A$  at a distance  $z$ , the flux of neutrinos is given by

$$\phi_\nu = \left( \frac{2\gamma}{1 + \gamma^2\theta^2} \right)^2 \frac{A}{4\pi z^2} \quad (2.2)$$

These distributions are plotted for NOvA's baseline of 810 km in Figure 2.6. As we move further off-axis, the energy distribution becomes more narrowly peaked at lower energies, and at 14 mrad it lines up with the first oscillation maximum for  $\nu_\mu \rightarrow \nu_e$  oscillations for a 810 km baseline. Having the peak of our energy distribution here maximizes the amount of  $\nu_e$  appearance events we can observe. Alongside the increased flux in the region of interest, the narrow peak in energy results in reduced backgrounds from neutral current (NC) neutrino interactions, allowing even better measurement of  $\nu_e$  appearance signals. Because no charged lepton is produced in a neutral current event, the neutrino often carries

away most of the energy, and so high-energy NC events will “feed down” to lower energy and may look like  $\nu_e$  events. Since we have a narrow-band beam, there are much fewer high-energy events to feed down into our signal region, whereas a wide-band beam would lead to larger NC backgrounds [38].



(a) Neutrino energy as a function of parent pion energy (Eq. 2.1) for different off-axis placements of the FD.

(b) Top: Event rates as a function of energy for different off-axis placements of the FD assuming an 810 km baseline. Bottom:  $\nu_\mu \rightarrow \nu_e$  oscillation probability as a function of energy. Note the first oscillation maximum is close to the distribution at 14 mrad.

**Figure 2.6:** Neutrino energy (a) and event rates (b) of neutrinos at the FD for different off-axis angles, assuming a baseline of 810 km. Taken from [39] and [40].

## 2.2.6 Final Beam Components

Beam Mode	$\nu_\mu$	$\bar{\nu}_\mu$	$\nu_e + \bar{\nu}_e$
FHC	95%	4%	1%
RHC	6%	93%	1%

**Table 2.1:** Summary of the predicted flavor profile of the NuMI beam in each running mode. Numbers represent flux integrated over the 1-5 GeV energy range in the absence of oscillations.

The calculated amount of each neutrino flavor in the final beam is summarized in Table 2.1. The beam is composed primarily of muon neutrinos in FHC mode, and muon antineutrinos in RHC mode, with a small amount of wrong-sign and  $\nu_e$  ( $\bar{\nu}_e$ ) backgrounds in each case. The primary decay modes for the pions and kaons that produce the muon (anti)neutrinos are summarized in Table 2.2.

Decay Mode	Branching Fraction (%)
$\pi^+ \rightarrow \nu_\mu + \mu^+$	$99.98770 \pm 0.00004$
$\pi^- \rightarrow \bar{\nu}_\mu + \mu^-$	$99.98770 \pm 0.00004$
$K^+ \rightarrow \nu_\mu + \mu^+$	$63.56 \pm 0.11$
$K^- \rightarrow \bar{\nu}_\mu + \mu^-$	$63.56 \pm 0.11$

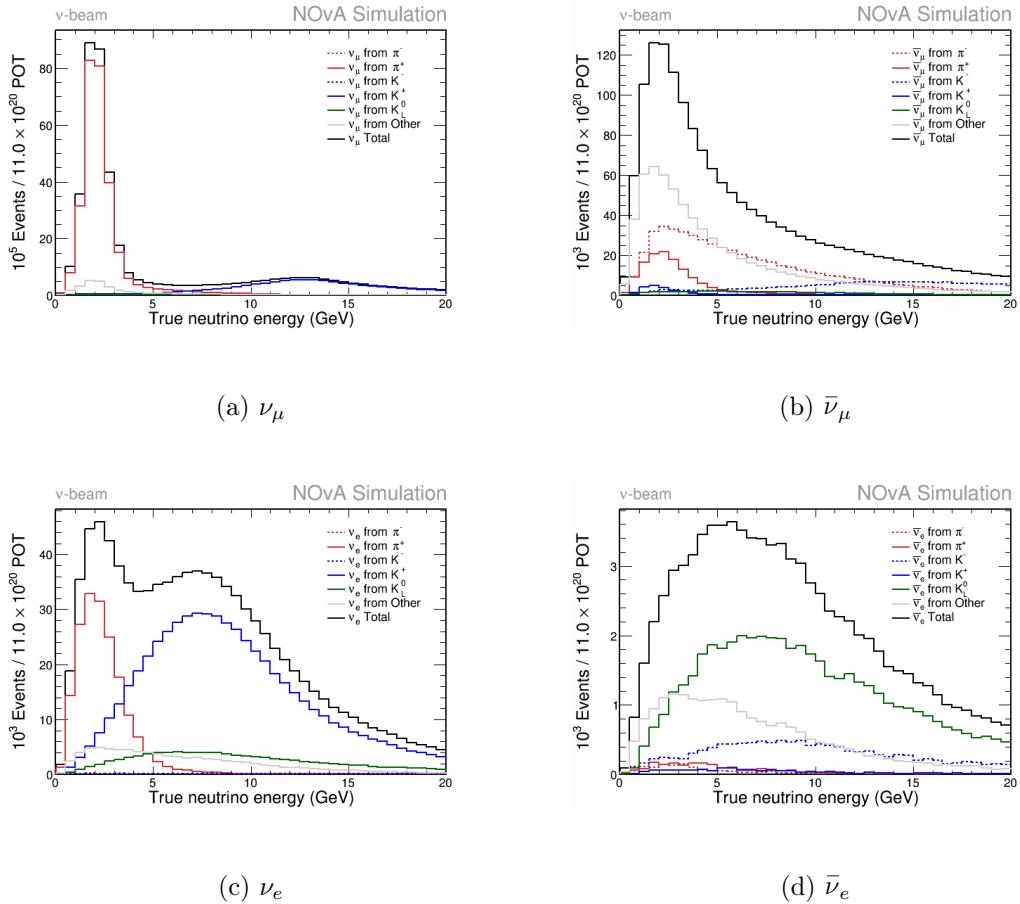
**Table 2.2:** The primary decay modes of the mesons that produce the neutrinos for our beam, along with their corresponding decay fractions [23].

Alongside the wrong-sign neutrinos discussed in Sec. 2.2.3, the other significant intrinsic beam background is a small amount of electron neutrino contamination coming from muon decay, as well as from some of the rarer kaon decay processes. These are summarized in Table 2.3

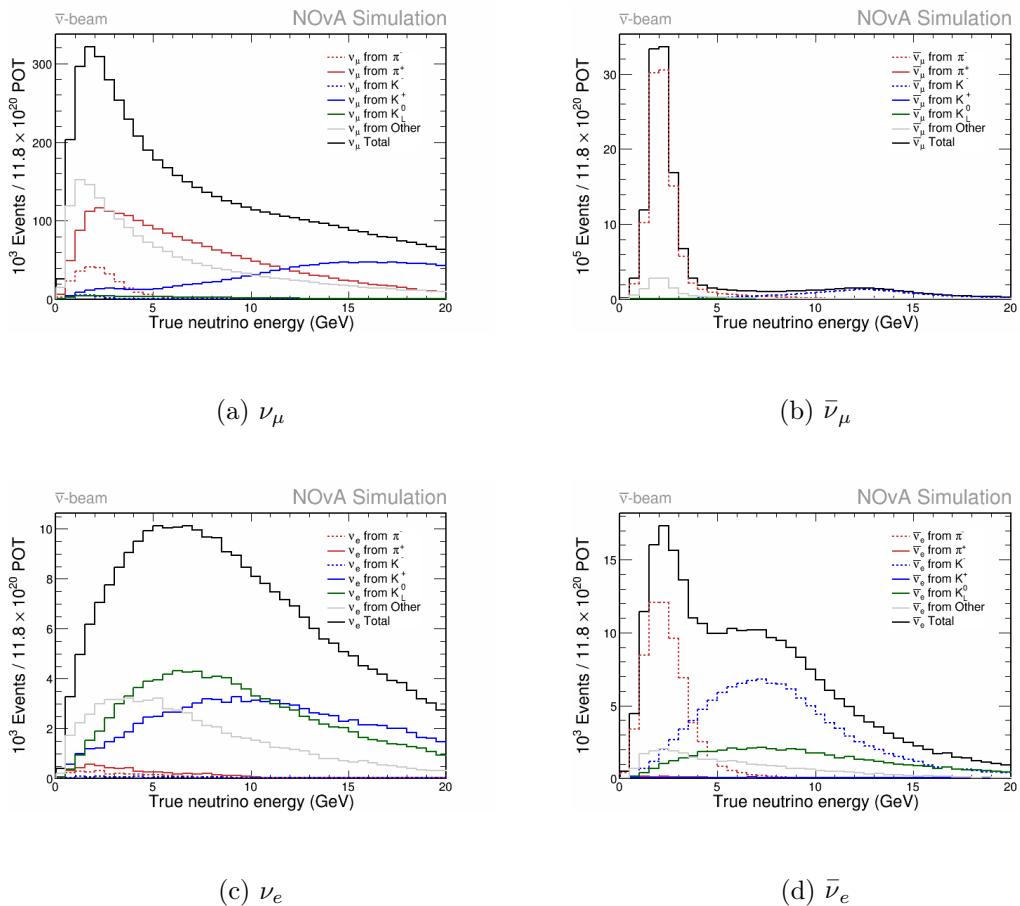
Decay Mode	Branching Fraction (%)
$\mu^+ \rightarrow \bar{\nu}_\mu + e^+ + \nu_e$	$\approx 100$
$\mu^- \rightarrow \nu_\mu + e^- + \bar{\nu}_e$	$\approx 100$
$K^+ \rightarrow \pi^0 + e^+ + \nu_e$	$5.07 \pm 0.04$
$K^- \rightarrow \pi^0 + e^- + \bar{\nu}_e$	$5.07 \pm 0.04$

**Table 2.3:** The decay modes of particles that contribute the most to the intrinsic electron neutrino backgrounds in our beam, along with their corresponding decay fractions [23].

Figures 2.7 and 2.8 show the beam flux distributions at the ND for each of the different neutrino types, broken down by the parent particle that decayed to produce it. The pion decays contribute the most around the desired beam energy range of 2 GeV, while the kaons mostly contribute to the higher energy neutrinos around 7 GeV for  $\nu_e$  ( $\bar{\nu}_e$ ) and 12 GeV for  $\nu_\mu$  ( $\bar{\nu}_\mu$ ).



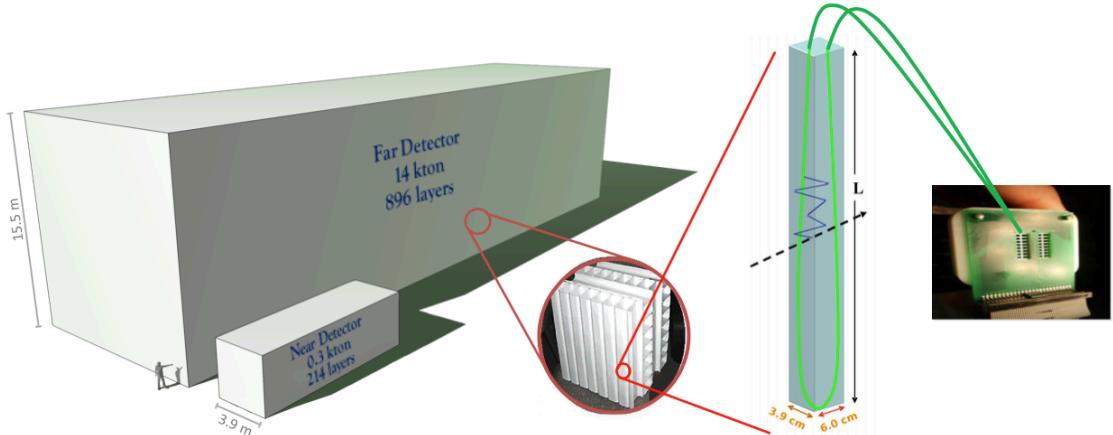
**Figure 2.7:** FHC energy distributions for different neutrino types at the ND, broken down by parent meson. Plots taken from [41]



**Figure 2.8:** RHC energy distributions for different neutrino flavors at the ND, broken down by parent meson. Plots taken from [41]

## 2.3 The NOvA Detectors

NOvA uses two low-Z segmented tracking calorimeter detectors to measure neutrinos from the NuMI beam before and after oscillations. The two detectors are made to be functionally identical, which will allow some cancellation of systematic uncertainties related to beam composition and cross sections (discussed further in Sec. 5.5). The detectors differ mainly in their size (shown in Fig. 2.9) and placement relative to the Earth’s surface (recall Fig. 2.1). The 0.3 kiloton near detector has dimensions 3.9 m wide  $\times$  3.9 m tall  $\times$  15.9 m long. It is located 105 m underground at Fermilab, about 1 km from the beam source. The much-larger 14 kton far detector has dimensions 15 m wide  $\times$  15 m tall  $\times$  60 m long and is located on the surface 810 km away in Ash River Minnesota. Both detectors are centered slightly off the beam axis, with the FD at an angle of 14.6 mrad.



**Figure 2.9:** Diagram showing the scale of the detectors, along with a zoom-in on the essential components. On the right is the basic unit, the cell, with its wavelength shifting fiber connected to the pixel of an APD. These cells are stacked to form planes, which are then alternated to make up the full detector. From [40].

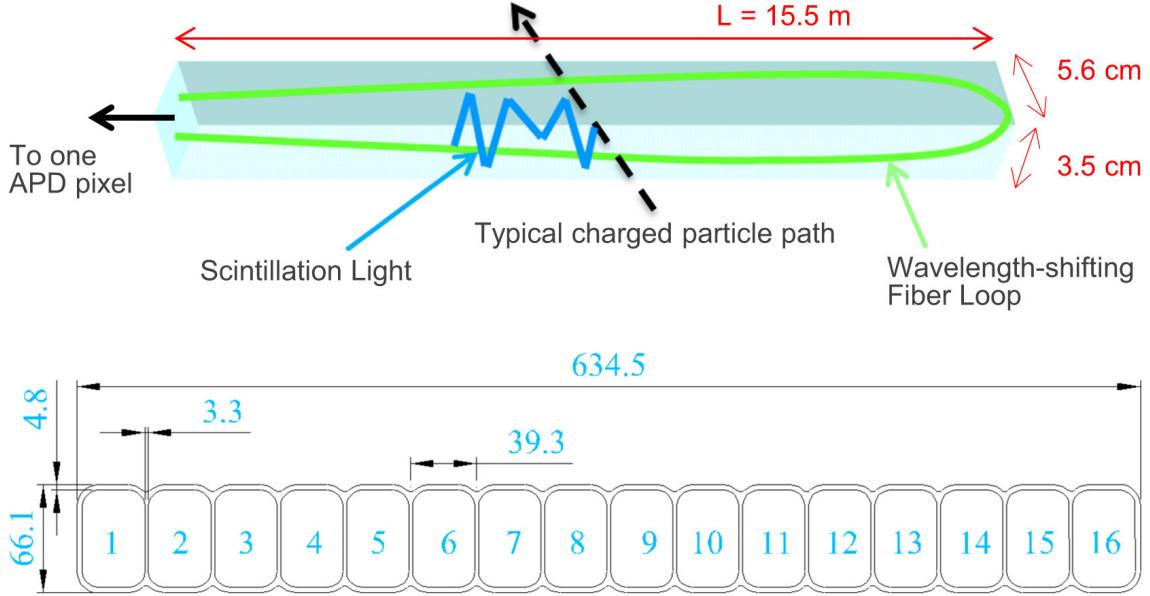
In order to meet our physics goals of measuring the oscillation channels  $\nu_\mu \rightarrow \nu_\mu$  and  $\nu_\mu \rightarrow \nu_e$ , the detectors need to be able to identify and differentiate between particles that make tracks, like the muons produced in  $\nu_\mu$  CC events, and those that make showers such as the electrons in  $\nu_e$  CC events. To this end, low-Z materials were used which give

a long radiation length. In combination with a fine-grained segmented detector, we have the position resolution to be able to distinguish electron showers from muon tracks and other hadronic backgrounds. Of particular note is the background coming from neutral current (NC) neutrino events. In NC neutrino interactions a  $\pi^0$  can be produced, which then promptly decays to two photons. If the showers from these two photons coincide it can mimic the shower made by electrons in  $\nu_e$  events. However, the photons travel a short distance before showering, and with our detector's resolution we can see this as a displacement of the photon showers from the neutrino interaction vertex. The electron from a  $\nu_e$  event, on the other hand, will have its shower begin immediately at the interaction vertex. This section will go into more detail on how the detectors function in order to achieve this goal.

### 2.3.1 From Cells to Full Detector

The fundamental unit of the detector is a polyvinyl chloride (PVC) cell filled with liquid scintillator, with a loop of wavelength-shifting fiber-optic cable running the length of the cell to trap and transport light to the readout electronics. Figure 2.10 shows a diagram. The cell has a cross-section of  $6.6 \times 3.9$  cm, and a length of 15.5 m (3.9 m) in the FD (ND).

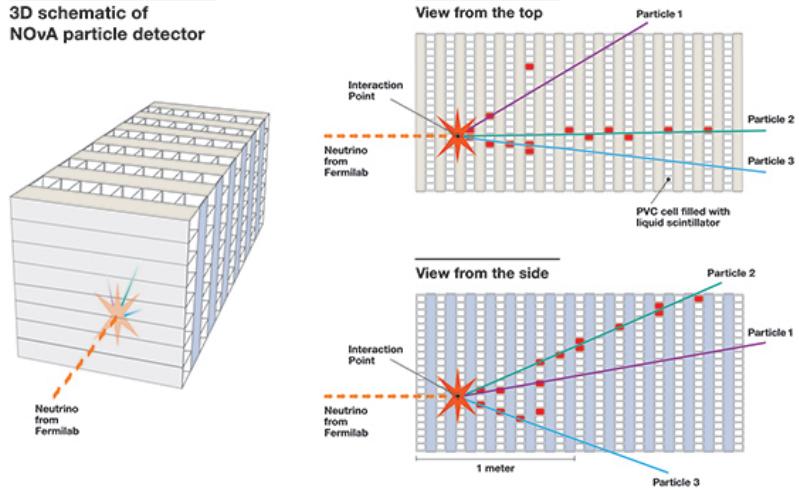
The PVC cells are made via an extrusion process where they are compounded with Titanium Dioxide to increase the reflectivity of the walls [42]. Simulations showed that, on average, light produced in the cell will be reflected 8 times before getting trapped and transported in the fiber, and so a 1% increase in reflectivity leads to  $\sim 8\%$  more light collected [42]. This is particularly relevant in the far detector, where the long length of the cells can lead to significant attenuation of light in the fibers. The scintillation liquid is made of 95% mineral oil, with the remaining 5% being pseudocumene and additional trace elements to increase the light yield. This liquid constitutes the 65% active mass of the detectors, while the PVC makes up the remaining 35% inactive material. The looped optical fiber has each end attached to one pixel of an avalanche photodiode (APD) at the



**Figure 2.10:** Diagram of essential cell components. Top dimensions refer to liquid volume. Images taken from [42].

end of the cell, which amplifies and digitizes the signal. This looped shape allows light to travel in either direction towards the APD, increasing light yield by a factor of 4 compared to a single non-looped fiber [38]. The fiber additionally shifts the collected light to the 490 - 550 nm range where the APDs have a higher quantum efficiency [43].

The cells are extruded in groups of 16 (see Fig. 2.10) and then stacked and glued together to form planes 384 (96) cells wide in the FD (ND). The planes are arranged in alternating horizontal and vertical orientations, which gives us both an “X view” and “Y view” of the detector (see Figure 2.11). The horizontal planes provide the Y view, and the vertical planes provide the X view. By combining both views, we can then get full 3D particle track reconstruction in our detector (more on how this works in Sec. 4.1). The FD has 896 total planes, and the ND has 214 planes. The final 22 planes at the downstream end of the ND are interleaved with 10cm-thick steel plates whose purpose is to range out any produced muons. This section of the ND is called the “muon catcher”. Its planes are slightly smaller than those in the main volume of the ND, standing 2/3rds as tall, since



**Figure 2.11:** Schematic showing what the X and Y views of the NOvA detector look like for an example neutrino event. Image from [44].

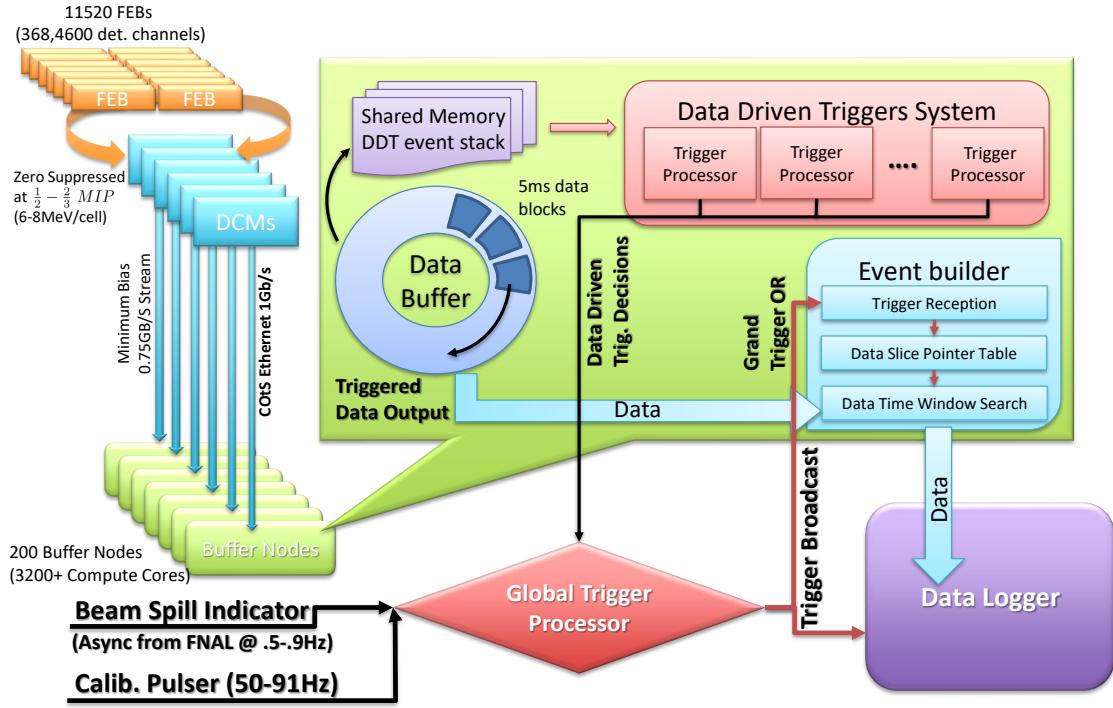
the steel was reused from a smaller prototype detector called the Near Detector On the Surface (NDOS).

### 2.3.2 The Data Acquisition System

Every cell constitutes one channel connected to one pixel of an APD. With 896 planes and 384 cells per plane that amounts to 344,064 channels in the FD (and 20,192 in the ND) whose signals need to be constantly read out and digitized. The Data Acquisition System (DAQ) handles all these signals simultaneously, while keeping the time synchronized across devices and triggering on relevant physics activity [45].

The APDs are the first step in the DAQ. Each of 32 pixels on an APD is a separate channel, reading in light from the fibers, amplifying them, and then converting it to an electric signal. To keep the electronic noise levels low, the APDs are connected to thermal-electric coolers which maintain a  $-15^{\circ}$  C operating temperature, while dry nitrogen gas is pumped into the APD housing to prevent condensation on the connecting surface.

Each APD constantly sends its 32 independent signals to a Front-End Board (FEB), which contains relevant electronics for low-level processing. An Application Specific In-



**Figure 2.12:** Overview of the Data Acquisition System for the FD. The ND is similar, but with a smaller number of channels. From [46].

Integrated Circuit (ASIC) first shapes the pulses to give them a characteristic rise and fall time. An Analogue to Digital Converter (ADC) then converts the shaped pulses to a digital signal. Finally, a Field Programmable Gate Array (FPGA) adds a timestamp to the signals, and passes any with an amplitude above a threshold to the Data Concentrator Modules (DCMs). Signals above threshold that get read out in this way are also referred to as "hits". Each DCM collects signals from 64 FEBs to consolidate into a  $50 \mu\text{s}$ -long data packet called a "microslice". These microslices are sent to a farm of buffer nodes where they are stored for up to 16 minutes while their information is analyzed and a triggering decision is made. Because of the large number of channels and high rate of cosmic events ( $\sim 150$  kHz) in the far detector, these microslices can amount to  $>1$  Gb/s of data being

created.

There are several different types of triggers that can cause data to be permanently saved by the Data Logger for later analysis. Timing-based triggers save all data around a specific point in time. For example, the FD cosmic minimum-bias trigger saves a  $550 \mu\text{s}$  window of data every 100 ms. Signal-based triggers save a specific amount of data when an external signal is received. An example of this is the NuMI trigger, which saves  $500 \mu\text{s}$  of data centered around the  $10 \mu\text{s}$  beam-spill. Finally, we have data-driven triggers that use information in the microslices to determine whether to trigger or not. The supernova trigger looks for a large number of small clusters of hits as an indicator that a supernova is occurring, and will read out a variable window of data based on the clusters [47].

In order to facilitate this precise triggering and hit readout, a strict and uniform timing system is required. Time Distribution Units (TDUs) connect to multiple DCMs and keep them all synced to the same local time provided by a Master TDU (MTDU) which is itself synced to a Global Positioning System (GPS) to maintain consistent timing between the two detectors and the Fermilab accelerator complex providing the beam.

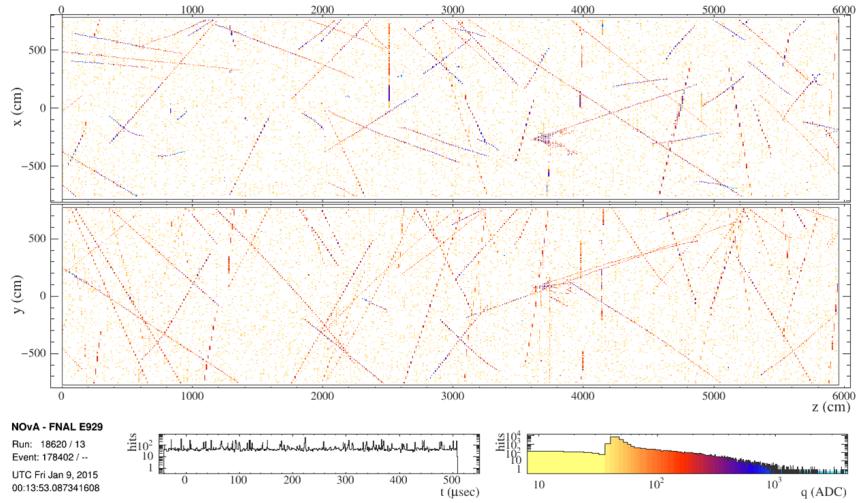
### 2.3.3 File organization

Data from each trigger is saved separately to its own set of files. Each file will contain many instances of that trigger, organized by time into runs and subruns. A run consists of multiple subruns, and the length of a subrun varies by detector and running conditions. In the near detector, a new subrun is started approximately every hour. At the FD, it is determined by the trigger rate. Each trigger is writing to its own set of files and when a total file size limit is reached for the current subrun, a new subrun is started. This is usually caused by the cosmic trigger, and results in a few subruns per hour. Groups of runs can be organized into data taking periods, sometimes called “epochs”, which are determined by larger changes in beam running conditions, such as switching from FHC to RHC running, or entering a beam shutdown period in the summer months.

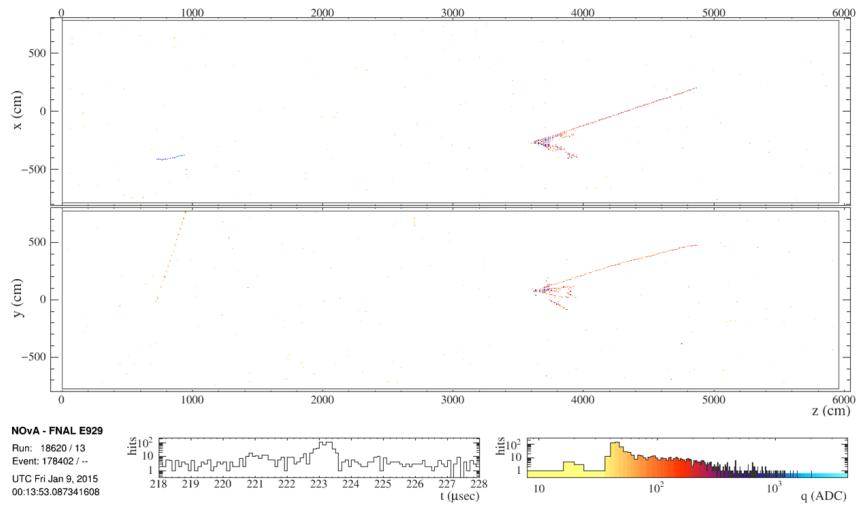
### 2.3.4 Event Display and Detector Differences

Figures 2.13 and 2.14 show event displays of NuMI trigger data for the FD and ND, which are useful for illustrating some of the key differences in the two detectors. The FD, being located on the surface of the earth, experiences a high rate of cosmic events, which can be seen in Fig. 2.13a. We can mitigate this background by zooming in on only the data that occurs in the beam spill window. This is shown in Fig. 2.13b, revealing a beam neutrino event. The ND, being located underground, experiences significantly less cosmics. However its proximity to the beam source means it sees a much higher flux of neutrinos, with several events occurring every spill (compared to  $\sim 1$  a week at the FD). Additionally, there are neutrino events that occur in the rock surrounding the detector, with the particles produced in the events (primarily muons) being another background at the ND. The multiple neutrino events and rock muons can be seen in Fig. 2.14. To deal with this pile-up of events, the ND samples each channel of the APD at a higher rate than at the FD. This is shown in Fig. 2.14b where we can distinguish multiple tracks within a single beam spill. The rate at the ND is 8 MHz, or one digitization (ADC readout) every 125 ns. The rate at the FD is 2 MHz, or one digitization every 500 ns [48]. Beyond the lack of pile-up, the coarser readout timing at the FD is necessary since the large number of detector channels would lead to unmanageable file sizes from all the extra data.

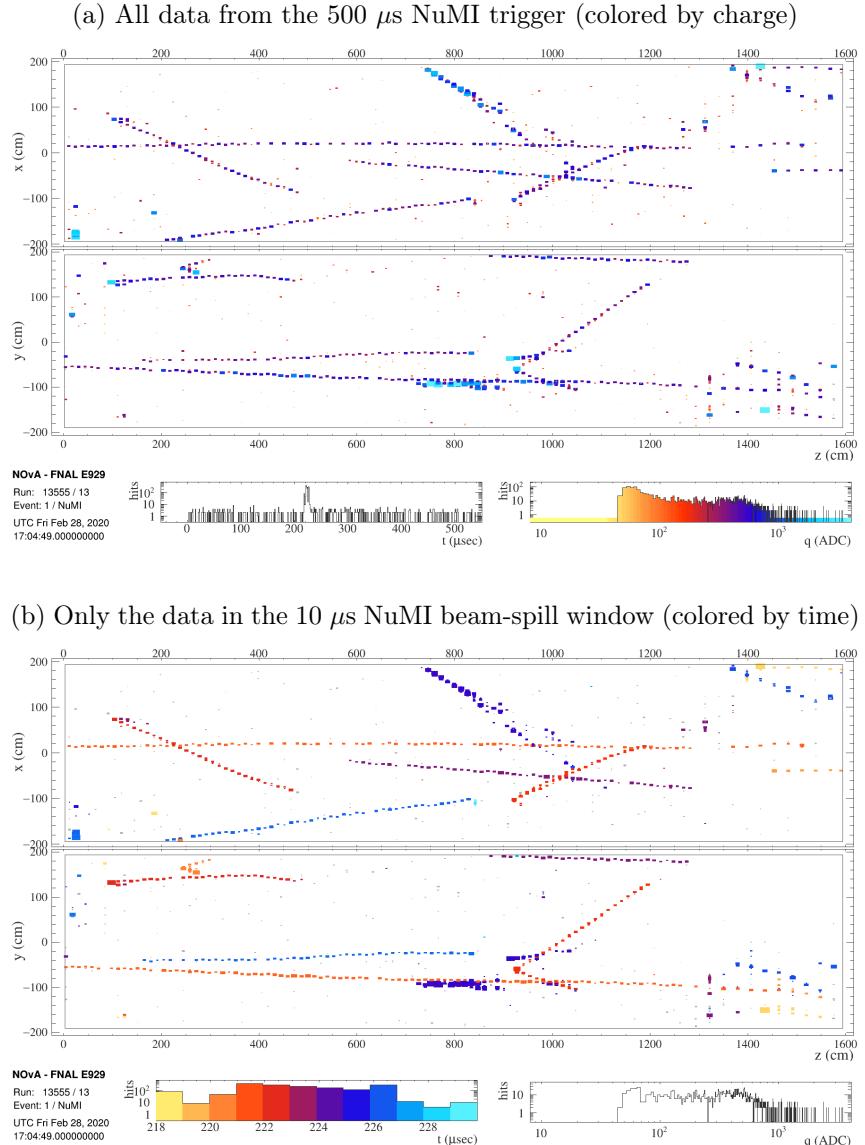
(a) All data from the  $500 \mu\text{s}$  NuMI trigger



(b) Only the data in the  $10 \mu\text{s}$  NuMI beam-spill window



**Figure 2.13:** A far detector event display, showing the high rate of cosmic events, as well as the effectiveness of isolating the beam-spill window. (a) shows all data from the NuMI trigger with the color indicating the charge collected by the APD, corresponding to the amount of energy deposited. (b) shows a zoom-in on the beam-spill window, revealing a muon neutrino event from the beam.



**Figure 2.14:** A near detector event display, showing multiple neutrino interactions as well as rock events occurring in a single beam spill. (a) shows all data from the NuMI trigger with the color indicating the charge collected by the APD, corresponding to the amount of energy deposited. (b) shows a zoom-in on the beam-spill window, where the color now indicates the time of the hit, to show off the precise timing resolution at the ND which allows us to deal with the event pile-up.

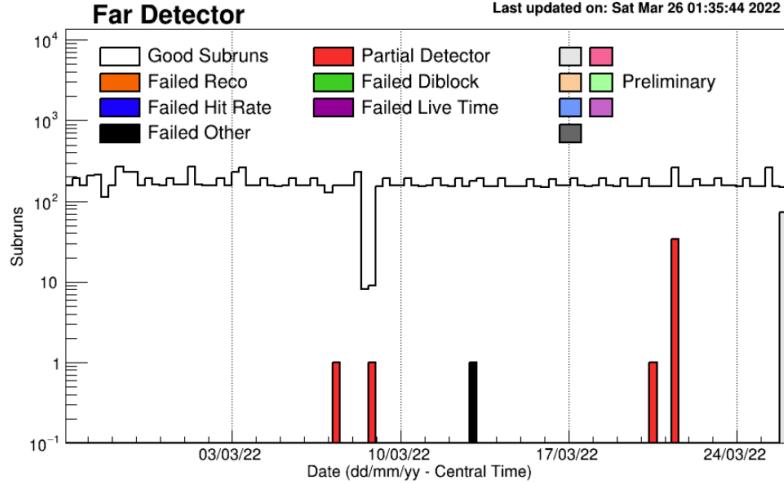
### 2.3.5 Data Quality Monitoring

NOvA maintains automated data quality monitoring scripts to ensure that the data being collected is of high quality and usable for our analyses. Both the beam and detectors are monitored, and runs of data that are taken when both are fully operational are labeled as “Good”. If either the detector or beam is in a full or partial failure mode, the run of data can be labeled as “Bad” and easily excluded from analyses (details on the cuts used in determining “Good” and “Bad” for this analysis will be discussed in Chapter 5). At the FD, in addition to labeling a full run as bad, we can also apply a “mask” to only exclude the regions of the detector that are undergoing a failure mode, and keep the data for the rest of the detector. This can extend down to labeling individual channels as good/bad. Figure 2.16 shows an example of one of the plots that are used to monitor the detectors, the average hit rate of the FEBs. Figure 2.15 shows the Good Runs monitoring, with the various failure modes labeled. A description of each mode is given below [49].

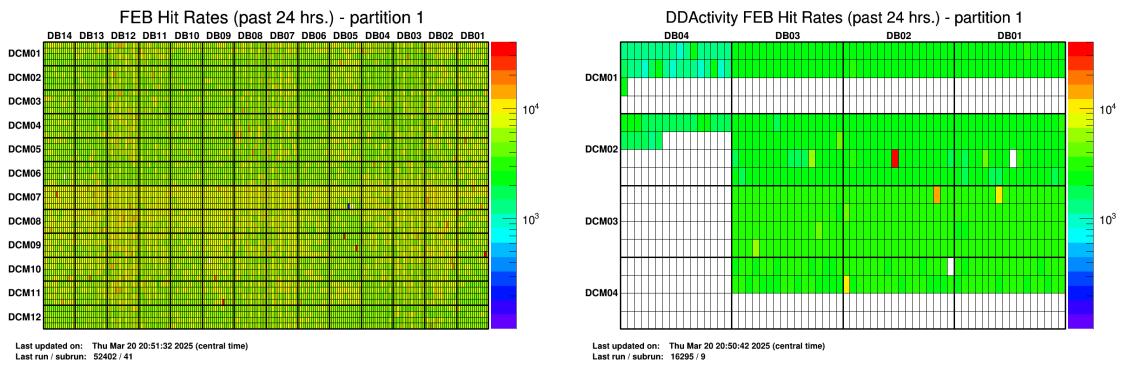
- Failed Reco - Too many or too few tracks; generally occurs when one or more DCM(s) is down.
- Failed Hit Rate - Median hit rate too high/low; beam instability is the main cause (ND).
- Failed Diblock<sup>1</sup>/Partial Detector - Part of the detector missing or a part of the detector having too high/low hit rate; general causes are detector crash, thunderstorm, issues with the diblock, etc.
- Failed Other - Incorrect timestamp or no activity; general cause is a change in Day-light Saving Time.
- Failed Livetime - Fewer than 1 k triggers in run; low POT (ND) .

---

<sup>1</sup>A Diblock is a group of 64 contiguous planes.



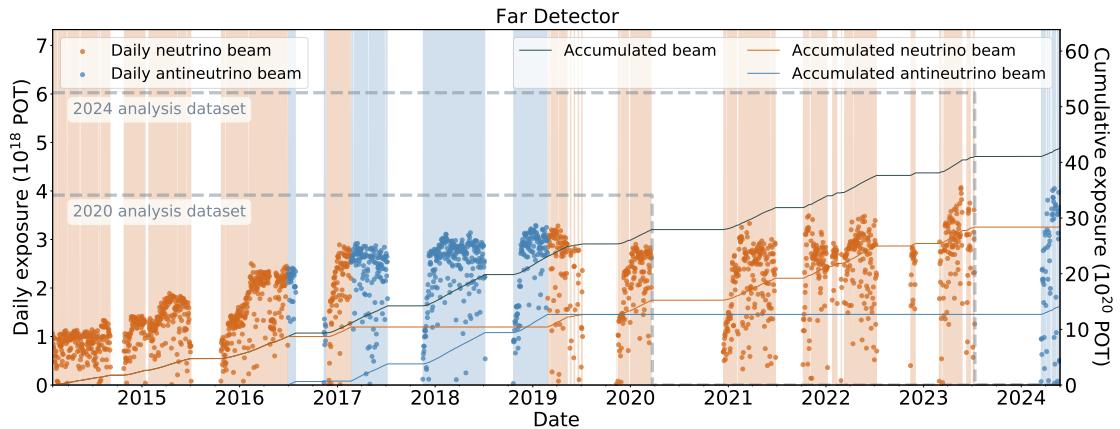
**Figure 2.15:** A plot of the Good Runs monitor, which indicates whether recent data collected is labeled as “Good” or is undergoing some type of failure mode. It is continually updated, with new data populated on the right. A description of the various modes is listed in the text. From [49].



**Figure 2.16:** Online Monitoring plots showing the FEB hit rate over a 24 hour period at the FD (left) and ND (right). The hit rate is one variable used in determining whether individual channels or runs of data are labeled as “Good” or “Bad”. Channels with abnormally high or low hit rates would be labeled Bad and excluded.

### 2.3.6 Total Data Collected

Since we have no practical way to measure the exact number of neutrinos produced by our beam, we instead use the number of protons delivered to the target as our metric for the amount of data collected. The unit is referred to as protons on target, or POT. Figure 2.17 shows the accumulated beam exposure at the FD across NOvA’s 10 years of operation. In addition, individual data points show the daily POT accumulated, which you can see increase over time as the beam was upgraded to provide more POT per spill, and reduced time between spills. For this thesis, we analyze a total FD neutrino beam exposure of  $26.61 \times 10^{20}$  protons-on-target (POT) corresponding to 885.53 s of beam-on time. This is combined with an antineutrino beam exposure of  $12.50 \times 10^{20}$  POT over 322.59 s of beam-on time.



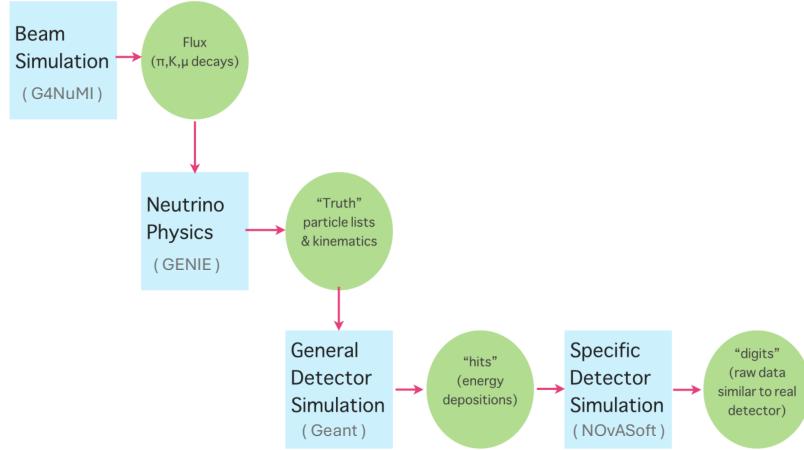
**Figure 2.17:** Accumulated beam exposure at the FD over time. Neutrino beam mode is indicated in orange, and antineutrino beam mode in blue. White regions are when the beam was off. This thesis uses the 2024 analysis dataset, while prior results used the 2020 analysis dataset. [50]

## Chapter 3

# Detector Simulation & Calibration

To interpret our data, we need an accurate simulation of the beam and detector, and the ability to convert the quantities we measure to meaningful physical units consistently across the detectors. This is the job of the simulation and calibration software.

### 3.1 Simulation



**Figure 3.1:** Flowchart of our simulation chain. Blue squares indicate the processes being simulated (with the software package in gray), and green circles indicate data products produced at each stage and used as input for the next. Modified from [51].

Figure 3.1 shows a flowchart of the simulation chain used in NOvA, and the data

products produced at each stage. A variety of Monte-Carlo simulation software is used, including GENIE [52], GEANT4 [53], and detailed models of the beamline and detector geometries. To complement the base simulation, we make data-driven corrections at several stages in order to improve our data/MC agreement.

### 3.1.1 Beam Simulation

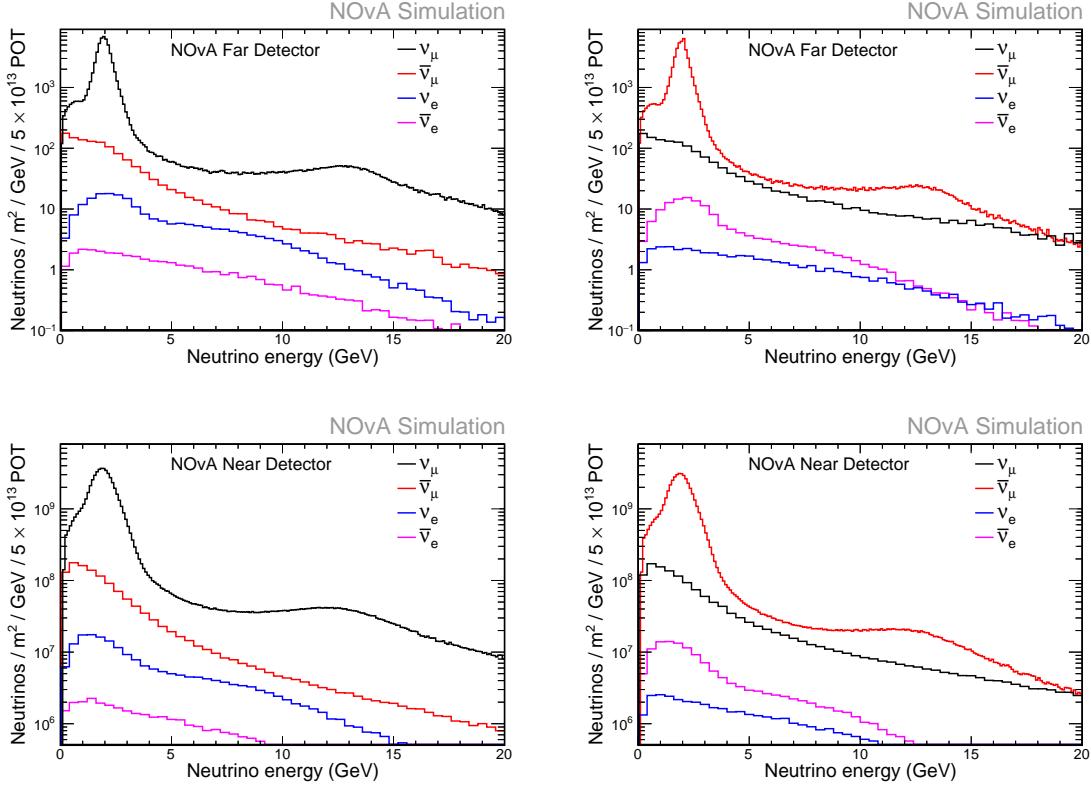
The goal of the beam simulation is to produce a set of flux files which describe the energy, flavor, and direction of neutrinos produced, along with kinematic information about the ancestor particles that lead to their production. It takes as input a simulated 120 GeV beam of protons with Gaussian position and angular distributions. It then uses a detailed model of the beamline geometry described in Sec. 2.2, called G4NuMI, along with GEANT4 to simulate the production of hadrons in the NuMI target, followed by their transport, focusing, and decay to neutrinos. The decay of the hadron and production of a neutrino ends this stage of the simulation.

The Package to Predict the Flux (PPFX) [55] is used to correct the baseline simulation, using constraints from hadron production experiments including NA49 [56]. It was developed for the NuMI beam within the MINERVA collaboration, and works by applying weights to individual neutrino events based on the kinematic information and interaction history of the hadrons that produced the neutrino. Figure 3.2 shows the predicted flux at the NOvA ND and FD with the PPFX corrections applied.

### 3.1.2 Detector Simulation

The next step after the beam simulation is the simulation of neutrino interactions in the detectors. At the neutrino energies that our beam produces, there is a variety of different interaction types that can occur. Figure 3.3 shows the  $\nu_\mu$  cross section on carbon as a function of energy for these different interaction types [16].

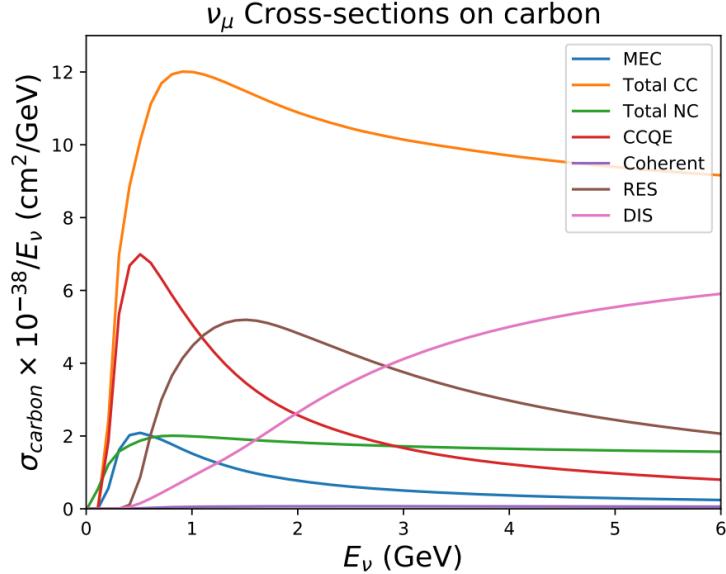
NOvA uses GENIE [52] to simulate the cross section and interaction of neutrinos in the detectors and surrounding material. GENIE uses the flux files from the beam



**Figure 3.2:** Neutrino flux with PPFX corrections applied at the FD (top) and ND (bottom) for FHC (left) and RHC (right) beam modes [54].

simulation, and a detector geometry file, combined with global neutrino cross section data and theoretical models to determine if a neutrino will interact along its predicted path. When an interaction occurs, GENIE will determine a vertex location and kinematic process for the interaction, and produce a set of final state particles that can be handed off to the next stage of the simulation [51].

In their default state, the GENIE models do not agree well with the NOvA data. As with the beam simulation, data from the NOvA ND and external measurements are used to form corrections to the simulation. These corrections are then applied at analysis time as weights to the individual simulated neutrino events. Table 3.1 lists the physics models used by GENIE for the different interaction types, along with the corresponding corrections that are applied. After first correcting the other interaction processes, we construct a NOvA-

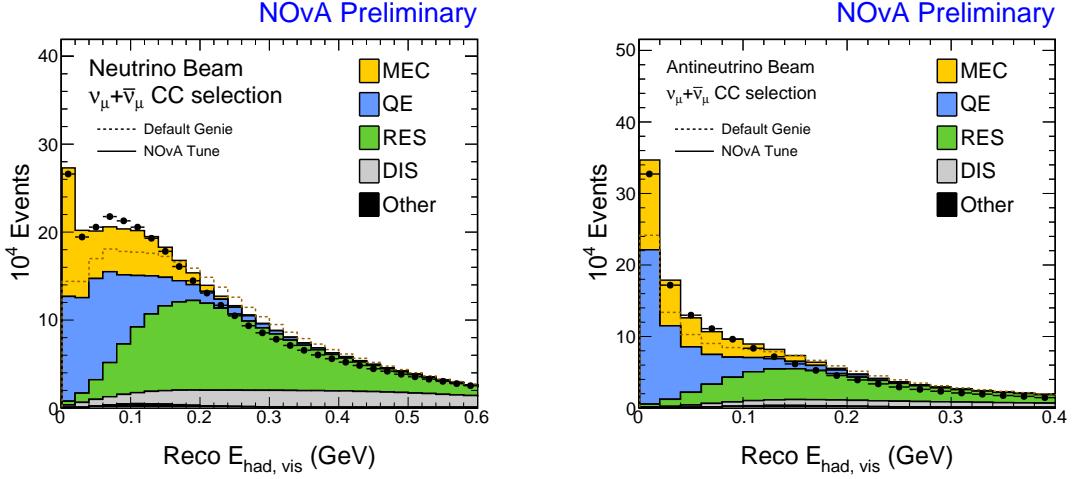


**Figure 3.3:** Muon neutrino cross section on carbon for different interaction processes, as a function of energy. From [16]

data-based correction to the Meson Exchange Current (MEC) model. This interaction describes charged-current neutrino scattering off of a correlated pair of nucleons via meson exchange. It contributes significantly to NOvA’s cross section, but is underestimated in the default configuration of GENIE [57]. The correction is constructed based on a fit to NOvA ND Data using reconstructed visible hadronic energy and three momentum transfer [58]. The effect of the complete cross-section weights, which combine the effects of all the corrections, can be seen in Fig. 3.4. The simulation agrees significantly better with the data after tuning.

The list of final state particles produced in the neutrino interactions by GENIE are then used as input to GEANT4, along with a detector geometry file. GEANT4 then simulates the propagation of those particles through the detectors, and any subsequent interactions. The final output of this stage of simulation is the location and size of the energy depositions made by each particle as it traverses the detector.

These energy depositions are then converted to number of photons  $N_\gamma$  using the fol-



**Figure 3.4:** Data/MC comparison of the reconstructed visible hadronic energy of  $\nu_\mu$  CC selected events in the ND, for FHC (left) and RHC (right) beam modes. The black points are the data and the simulation with cross-section weights applied is shown as a stacked histogram with the different interaction modes indicated. The dashed line is the default GENIE simulation before any tuning. From [59]

lowing equation:

$$N_\gamma = F_{\text{view}}(Y_s E_B + \epsilon_C C_\gamma) , \quad (3.1)$$

where  $F_{\text{view}}$  are overall scaling factors for the X and Y views in each detector (four in total),  $Y_s$  is an overall scaling for scintillation energy,  $\epsilon_C$  is an overall scaling factor for Cherenkov energy,  $C_\gamma$  is energy deposited due to Cherenkov photons, and  $E_B$  is the Birks energy, which is determined from Birks Law

$$\frac{dL}{dr} = S \frac{\frac{dE}{dr}}{1 + k_B \frac{dE}{dr}} \quad (3.2)$$

where  $\frac{dL}{dr}$  is the light emitted per pathlength,  $\frac{dE}{dr}$  is the energy loss per pathlength,  $S$  is the scintillator efficiency (set to 1 after calibration is applied) and  $k_B$  is the Birks constant. This equation is used to modify the energy distribution to include the effect of quenching, which is any loss in light yield of the scintillator due to inherent material properties, such as when photons are reabsorbed by the scintillation material. This occurs when the  $\frac{dE}{dx}$  is

CC Interaction Process	Model	Correction
Initial State	València Local Fermi Gas Model [60]	N/A
QE	València + Z-Expansion [60]	Random Phase Approximation (RPA) included to model effect of long-range nuclear correlations [60]
MEC	València [61, 62]	Parameters tuned to match NOvA ND Data
Resonant Production	Berger-Seghal [63]	Parameters tuned to match neutrino-deuteron scattering data [64]
DIS	Bodek-Yang [65]	Parameters tuned to match neutrino-deuteron scattering data [64]
FSI Model	GENIE hN intranuclear cascade model [66]	Parameters tuned to match pion-nucleus scattering data [67–73]

**Table 3.1:** Physics models in the GENIE version used for this analysis, along with the corrections applied to each model to improve data/MC agreement.

very high, so it mostly affects the ends of tracks. This effect is parameterized by the Birks constant  $k_B$ , which has been measured by NOvA to be  $0.01155 \pm 0.00065$  g/(MeV cm<sup>2</sup>) for our scintillator [74]. We once again apply a data-driven correction to our model, with the scaling factors  $F_{\text{view}}$ ,  $Y_S$ , and  $\epsilon_C$  determined from a least-squares fit to data. Four different sets of data and MC samples are used to perform the joint fitting on the light level parameters. These samples are: ND muons, ND protons, ND cosmics, and FD cosmics. This light model tuning is also used to set the size of our systematic uncertainty on the light model, and is discussed further in Sec. 5.5.

The next step is to simulate the transport of light through the cell and to the APD via the fiber. While GEANT4 is capable of simulating this process through ray tracing, it is computationally expensive to do so for every photon. Instead, NOvA uses a custom light simulation software that uses templates to parameterize the collection of photons by the fiber, the transport of those photons up the fiber, and the response of the APD [75]. The initial position and number of photons is used as input, and values are drawn

from the templates for each process to determine the photo-electrons as a function of time for a given energy deposit. This step also takes into account attenuation along the fiber, using quality control measurements of fibers taken prior to detector construction [75]. To account for APD noise, the number of photo-electrons are smeared according to a log-normal distribution, with the final response matching the theoretical distribution well [76].

The final stage of the simulation models the pulse shaping and triggering logic of the DAQ readout electronics. The pulse shaping is modeled by a double exponential, one for the rise time and one for the fall time, that match the times of the ASIC on the FEB [75]. A dual-correlated sampling (DCS) method is used to determine if a hit is read out. The DCS value is given by

$$DCS_i = ADC_i - ADC_{i-3} \quad (3.3)$$

where  $ADC_i$  is the value of the current digitization sample, and  $ADC_{i-3}$  is the sample from 3 digitizations prior (as a reminder, the FD cells produce a digitization every 500 ns, and the ND every 125 ns). If the DCS value is above threshold, the hit is read out with the time and ADC of the hit recorded. The thresholds for each individual cell are pulled from a database to match with the corresponding cell in the real detectors. After triggering and reading out a hit, the APD will not re-trigger for a specified amount of time, called the dead-time.<sup>1</sup>

The end product of this simulation chain is a file containing detector hits that match the format seen in the data, but with the addition of truth information from the simulation. The truth information includes the particle types and interactions that lead to the hits, the true energy deposited, and kinematic and ancestry information for all particles. We can then pass both data and MC files through the same reconstruction chain to look for neutrino events and make inferences about the data. This is discussed further in the next

---

<sup>1</sup>The re-triggering logic is known to be mis-modeled in the simulation, leading to a data/MC discrepancy in APD deadtime. While this does not significantly affect the majority of signals, it has a large effect on delayed signals such as those from Michel electrons. This discrepancy and a future solution are discussed further in Sec. 6.4.

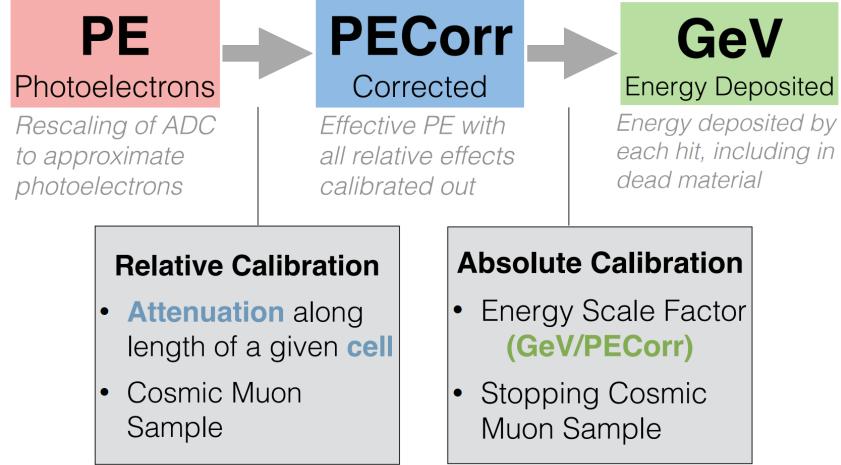
chapter.

### 3.1.3 Cosmic and Rock Events

Cosmic ray events at the FD, and rock muon events at the ND are both large sources of background that are computationally expensive to simulate since it requires modeling a much larger volume to get enough events that actually pass through the detectors. To solve this problem at the FD we instead overlay real data from our dedicated 10 Hz cosmic trigger onto the MC files to give a realistic cosmic background. At the ND, we simulate a smaller sample of neutrino events occurring in the rock in front of the detector, and overlay them onto the MC files, re-using events if necessary to match the POT/spill of a given run. The same is done for the FD, but the probability of a rock event occurring is much smaller due to the lower neutrino flux, so the FD rock events represent a small background in the analysis. For cases where a simulated cosmic sample is necessary, such as for calibration, we use the CRY generator [77].

## 3.2 Calibration

The calibration software allows us to convert our measured ADC values into an established unit of energy, which can be used to interpret physics results. The calibration is performed for both data and MC and done in two stages. The first stage, called the relative calibration, corrects for effects that cause a variation in the observed photoelectrons (PE) across individual detector elements, such as attenuation of light along the fibers. The output of this stage is a corrected photoelectron count for each hit, called PE<sub>Corr</sub>, which is uniform across the detector for a given energy deposition. The second stage, called the absolute calibration, finds a scaling factor to convert the PE<sub>Corr</sub> values to units of energy (GeV), which are useful to our physics analyses. A simple flowchart of the procedure is shown in Fig. 3.5.



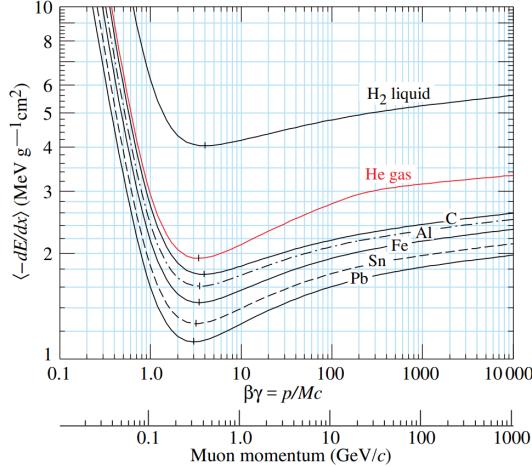
**Figure 3.5:** A simple flowchart of the calibration procedure. The top colored blocks show the units of hits at each stage. The gray arrows and blocks indicate the calibration procedure performed to take one data type to the next, along with its primary goal and sample used. Taken from [78]

Both stages of calibration are performed using cosmic muons since we have very large samples of cosmic data, ensuring high statistics in all cells, and the energy deposition of muons is well-understood using the Bethe-Bloch formula [79]:

$$\left\langle -\frac{dE}{dx} \right\rangle = K z^2 \frac{Z}{A} \frac{1}{\beta^2} \left[ \frac{1}{2} \ln \frac{2m_e c^2 \beta^2 \gamma^2 W_{\max}}{I^2} - \beta^2 - \frac{\delta(\gamma\beta)}{2} \right]. \quad (3.4)$$

Here, the constant  $K = 4\pi N_A r_e^2 m_e c^2$  combines Avogadro's number,  $N_A$ , the classical electron radius,  $r_e$ , and the rest energy of the electron,  $m_e c^2$ .  $W_{\max}$  is the maximum energy transfer to an electron in a single collision.  $Z/A$ ,  $I$  and  $\delta$  are the ratio of the atomic number to the mass number, mean excitation energy and the density effect correction to ionization energy loss for the target material, respectively. The charge, velocity and relativistic gamma factor of the incoming particle are given by  $ze$ ,  $\beta c$ , and  $\gamma$ , respectively.

The equation describes the mean energy loss rate,  $\frac{dE}{dx}$  of a heavy relativistic charged particle moving through a medium. The equation is valid in the region  $0.1 < \beta\gamma < 1000$  with an accuracy of a few percent [79]. The curve is plotted for muons incident on different media in Fig. 3.6.



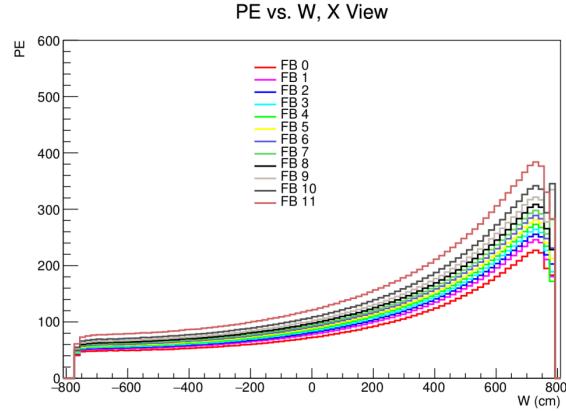
**Figure 3.6:** Bethe Bloch curve for muons in different materials. Taken from [79].

Particles located near the minimum of the Bethe-Bloch curve have a nearly constant  $dE/dx$  and are referred to as Minimum Ionizing Particles, or MIPs. For our detector, muons fall in the MIP region for the majority of their track, where they deposit approximately 1.8 MeV/cm [78]. This makes them ideal candidates for calibration. The relative calibration uses cosmic muons that traverse the entire detector, referred to as “through-going muons”. This ensures a roughly consistent energy deposition per length across the detector.

Absolute calibration uses cosmic muons that stop in the detector. Knowing the end position of the muon track, the Bethe-Bloch equation can be used to determine the exact energy deposited in a selected track region in GeV and then define a conversion factor between observed PEcorr and true energy. In each case, a simple track reconstruction algorithm called CosmicTrack is used to identify the cosmic muons, with separate selections applied afterward. The reconstruction is described in more detail in Sec. 4.1.

### 3.2.1 Relative Calibration

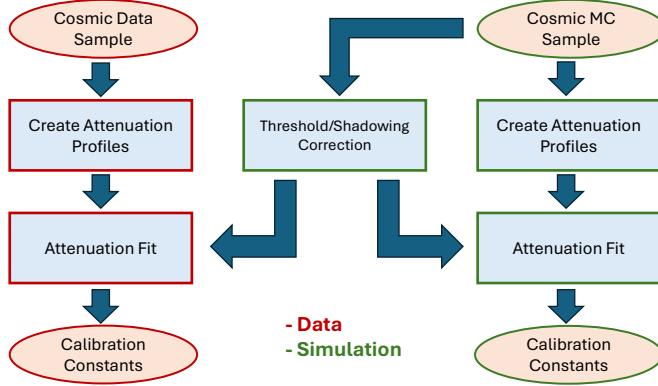
The purpose of the relative calibration is to correct for attenuation effects and cell-to-cell variations. Hits that occur far from the readout in a cell can have more light attenuated on the way to the APD and would thus appear to have a lower energy on average than hits



**Figure 3.7:** Simulated PE as a function of  $W$  (position along the cell) for FD cells in the X view, split by different fiber brightness bins. Higher  $W$  values are closer to the readout, with  $W = 0$  the center of the cell. The effect of attenuation can be seen as lower PE observed further from the readout.

close to the readout. This step also corrects for threshold and shadowing effects. Threshold effects are when the energy depositions occurring far from the detector readout may need to have a slight upward fluctuation in the number of photons produced in order to pass the hit threshold. This can bias the energy distribution to higher values far from the readout. Shadowing refers to self-shielding of the detector by its own mass, which causes the  $dE/dx$  to not be completely uniform throughout the detector. As a muon traverses the detector it will lose energy and move down the Bethe-Bloch curve, changing its  $dE/dx$ . These effects are most prominent at the FD, where the detector is larger and cells are much longer. The final effect covered by the relative calibration is differences in fiber brightness. This refers to the observation that different WLS fibers attenuate light at different rates (shown in Fig. 3.7). During detector construction, the attenuation of each fiber was measured, and the resulting distribution was split into 12 quantiles. The boundaries between groups are used as brightness ratings, ranging from 0-11, and each cell is assigned a brightness rating, and simulated with that attenuation value. The calibration is done separately for each fiber brightness bin to account for this.

Fig 3.8 shows a block diagram of the relative calibration process. We first create “attenuation profiles” separately for each cell in data and MC. These are 2D histograms of



**Figure 3.8:** Block diagram illustrating the main steps in the relative calibration.

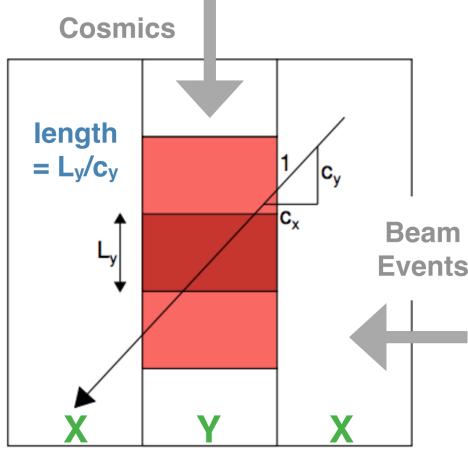
PE/cm vs. W (the position along the cell)<sup>2</sup>, which we then profile to get a 1D distribution of mean PE/cm vs. W for each cell.

To fill these histograms, we need to select appropriate hits from the throughgoing cosmic muon tracks. To calculate PE, we apply a simple scaling to the ADC of each hit<sup>3</sup>, but in order to calculate PE/cm we also need an accurate estimate of the pathlength of the muon through the cell. To ensure a proper pathlength estimate, we only consider so-called “tricell” hits which are located between two other hits in adjacent cells. Figure 3.9 illustrates this configuration. This guarantees the muon passed through opposite walls of the cell, and the pathlength can be found by taking the width of the cell and dividing it by the cosine of the track angle through the cell. The W of the hit is calculated by finding the nearest hits in the same track in the opposite view, and making a straight line between them.

Once the attenuation profile histograms are made, the next step is to calculate the threshold and shadowing corrections. This is done using only the MC tricell hits, since it requires truth information about the hit, with the same correction then applied to both

<sup>2</sup>W is a coordinate independent of the X,Y,Z coordinates normally used when referring to detector position. The same system is used for both horizontal and vertical cells, with W = 0 defined as the midpoint of the cell, more positive W being closer to the readout electronics, and more negative W being further from the readout.

<sup>3</sup>This scaling is different for each detector and depends on the gain of the APD.



**Figure 3.9:** Example of the tricell hit condition used in calibration for a horizontal (Y-view) cell. The dark red cell is selected due to having hits in the adjacent light red cells. The general direction of cosmic and beam neutrino events is indicated. Taken from [78]

data and MC. Instead of being done independently for every cell, we collate all planes in the same view, and the correction is calculated as a function of view, cell number within a plane, and fiber brightness. The correction is of the form

$$T = \frac{PE}{\lambda} * \frac{E_{true}}{E_{MIP}} , \quad (3.5)$$

where  $T$  is the combined “threshold and shielding” correction factor,  $PE$  is the simulated photoelectrons recorded at the readout,  $\lambda$  is the number of simulated photons which would be seen at the readout out in the absence of fluctuations,  $E_{true}$  is the true energy deposited in the cell and  $E_{MIP}$  is the energy you would expect to be deposited by a MIP based on the pathlength. The first factor corrects for the threshold effects from fluctuations in the number of photons produced, and the second factor corrects for variations in the MIP energy from self-shielding.

After applying this correction to the attenuation profiles, we fit the corrected hits to

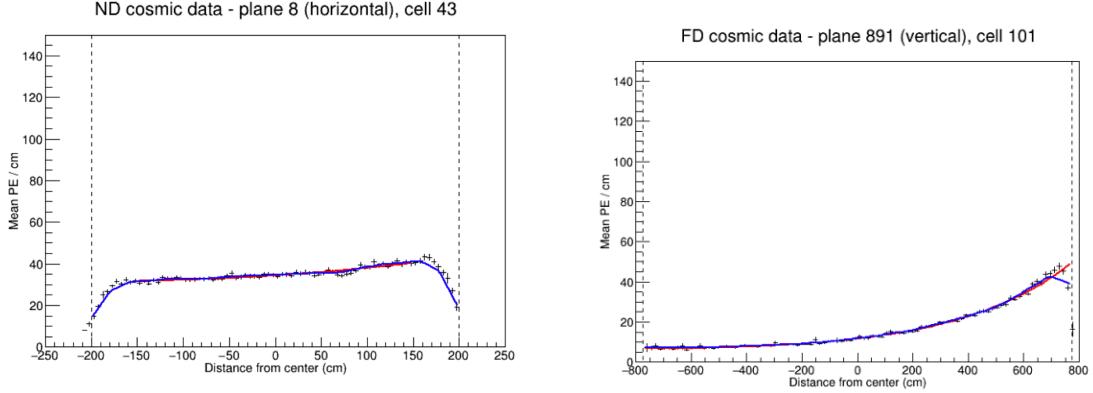
extract calibration constants. The fit is of the form

$$y = C + A \left( \exp \left( \frac{W}{X} \right) + \exp \left( -\frac{L+W}{X} \right) \right) \quad (3.6)$$

where  $y$  is the response,  $L$  is the cell length,  $W$  is the location of the hit in the cell, and  $C$ ,  $A$ , and  $X$  are free parameters in the fit, with  $X$  representing the attenuation length. The two exponentials represent the two paths a photon can take through the looped WLS fiber to the APD, with the first accounting for the shorter direct path, and the second representing the photon taking the longer path down the cell and back around (see Fig. 2.10).

The fit is performed for every cell independently in data, while for MC we collate planes in the same view, and the fit is done by view, cell, and fiber brightness. We do this for two main reasons; first, our simulated cosmic sample is much smaller than the cosmic data sample, due to the large simulation time required. So, to get the required statistics for the fit we collate planes. The other reason is that the simulated detector is assumed to be uniform plane-to-plane. We do not make this assumption for the real detector. In the ND muon catcher, which is less instrumented than the main near detector, there are significantly fewer tracks with the angles required to produce tricell hits. Here, the fit is done only by view in both data and MC.

The results of the fit for a typical cell are shown in Fig. 3.10. The fit performs well in the center regions of the cell, but does not describe well the ends of the cell, where a “roll-off” in the PE/cm is observed (see Fig. 3.10). To fix this, an additional locally weighted scatter plot smoothing (LOWESS) fit is done on the residuals of the first fit and used to correct the roll-offs near the ends of the cell. This is shown as the blue curve in Fig 3.10. The effect of the relative calibration applied to data and MC can be seen in Fig. 3.11.



**Figure 3.10:** The results of the attenuation fit for a typical cell in the ND (left) and FD (right). The red curve is the initial exponential fit, and the blue curve is the final fit with LOWESS correction.

### 3.2.2 Absolute Calibration

The last step in the calibration serves to find a conversion factor to take the PECorr values for each cell in each detector and express them in units of energy. This conversion factor is found by dividing the average true energy per pathlength (MeV/cm) from simulated muon hits by the average reconstructed energy/pathlength (PECorr/cm) of our reconstructed hits in both data and simulation.

To do this, we first find the appropriate muon hits used to calculate these mean values. Unlike with relative calibration, we look for muons that stop in our detector. To select these, we choose tracks from CosmicTrack that have a reconstructed end point within the fiducial volume of the detector<sup>4</sup>, as well as at least 1 Michel electron (the electron produced from muon decay) associated with the track. This process is described in Chapter 4. By knowing the end point of the muon, we can select hits in the region of the track where the muon is a MIP. For NOvA, this is the region 100-200 cm from the end of the track (see Fig. 3.12). The  $dE/dx$  of simulated muons is found to be flat to within 1.8% in this region [81]. This range was chosen to minimize the effect of track reconstruction failures where hits are missed at the end of the track. We also apply the tricell condition to these hits to

<sup>4</sup>The fiducial volume is a smaller internal volume of the detector, defined in Sec. 5.2.

ensure an accurate pathlength estimate.

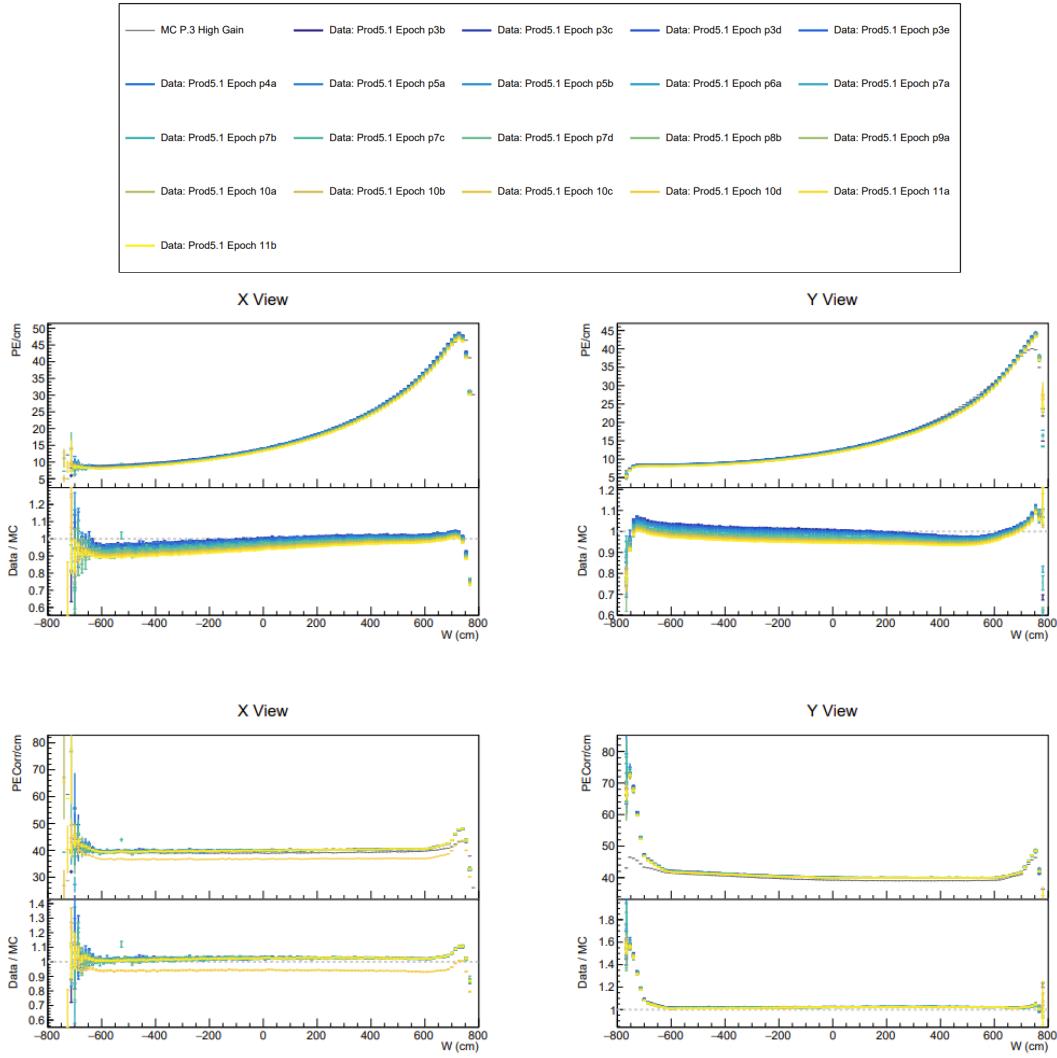
Once these hits are selected we fill a histogram of true MeV/cm from truth information in the MC hits, and histograms of PEcorr/cm from reconstructed hits in data and MC separately. We then find the mean of the histograms to be used for the correction. The conversion factor is then calculated as

$$\text{Correction} = \frac{\text{Average MeV/cm}}{\text{Average PEcorr/cm}} \quad (3.7)$$

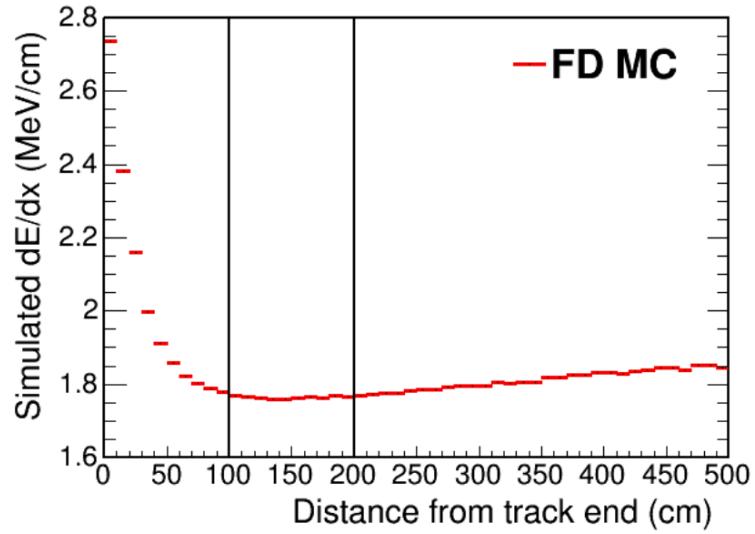
Separate factors are found for data and MC, and for each view. The separation in view is done to prevent potential bias coming from differences in the amount of hits between views. The cosmic muons are mostly downward-going, and will leave more hits in horizontal than in vertical cells. The final result of relative and absolute calibration can be seen summarized in Fig. 3.13.

### 3.2.3 Accounting For Drift

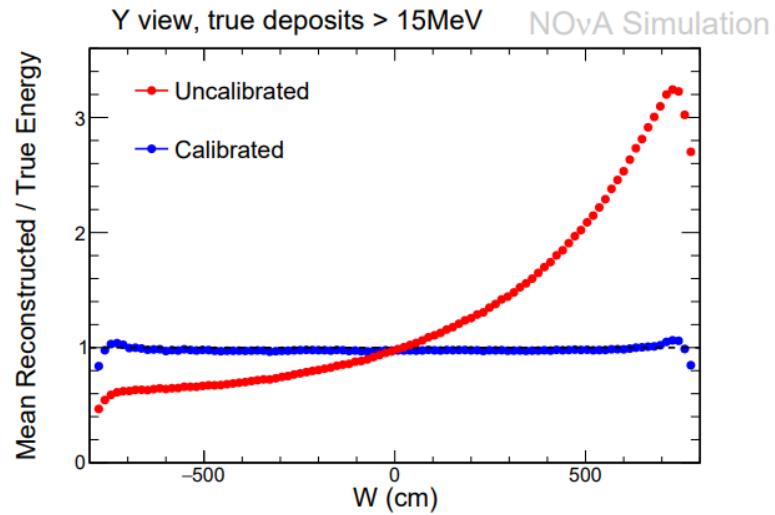
A decrease in the collected PE over time, known as “detector drift” has been observed in the FD data. It is expected to come from degradation over time of the scintillator material, fibers, and electronics. The drift can be observed in Fig. 3.14. To account for this, the calibration is done separately for different periods of time. For data taken before June 2021, it was split into “epochs”, which are run periods dictated by major changes to beam running conditions including changes in horn current, upgrades, or summer shutdown periods. For the majority of new data in this analysis, taken Sept 2021 - July 2023 (see Fig. 2.17) we have switched to a month-by-month calibration. No difference in effectiveness was observed using this new scheme [83].



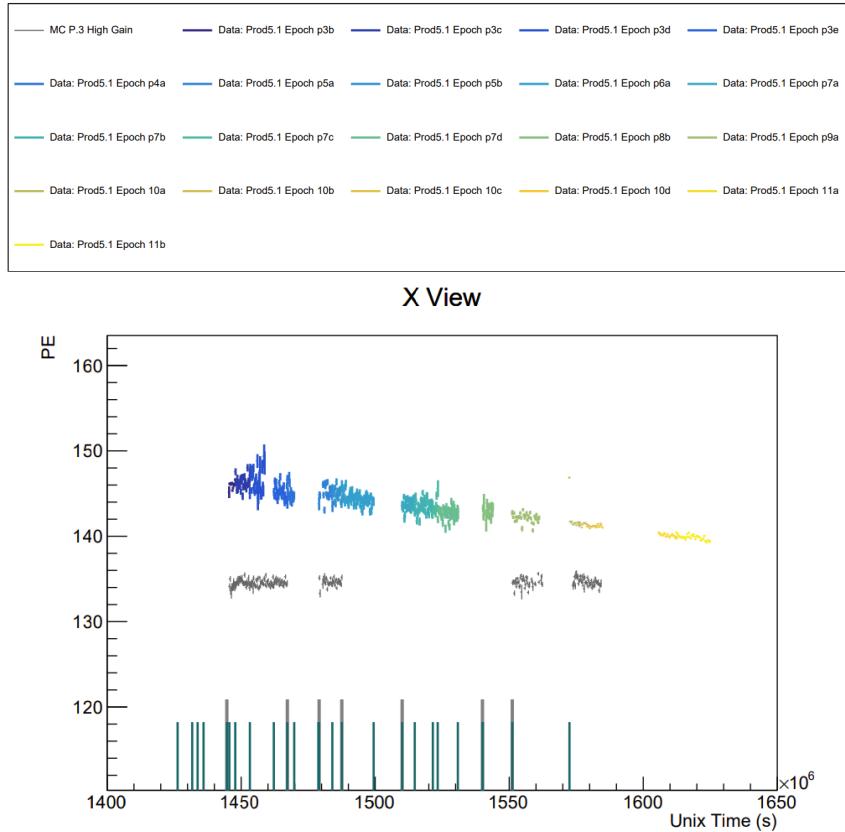
**Figure 3.11:** FD plots showing the effect of relative calibration, averaged across all cells, with  $\text{PE}/\text{cm}$  for each view on top and  $\text{PECorr}/\text{cm}$  for each view on the bottom. The different colored curves represent different data taking periods, with a comparison made to one MC period. These epochs represent data taken from October 2015 - June 2021. Epoch 10d has a significant deviation since it contained only one week of data. Constants from epoch 10c were used instead for that data, and it is only included here for completeness. Taken from [80].



**Figure 3.12:** Simulated muon  $dE/dx$ , with the selected track window used for calibration shown.



**Figure 3.13:** The mean reconstructed/true energy shown for simulated FD hits in the Y-view before/after calibration is applied. Relative calibration will flatten the curve, and absolute calibration will shift it up/down. Taken from [82].



**Figure 3.14:** Observed PE over time for the FD, averaged over all X-view cells. The steady decrease over time is known as detector drift. The colors correspond to different run periods. The vertical bars indicate the beginning of new epochs. The gray points are MC. [80]

## Chapter 4

# NOvA Software Infrastructure

Our analysis relies on our ability to reconstruct neutrino events from our data, and identify their flavor and energy. To do this we make use of a software chain that first builds up clusters of hits into events, finds topological features, and then feeds them into algorithms that can reliably estimate the flavor and energy.

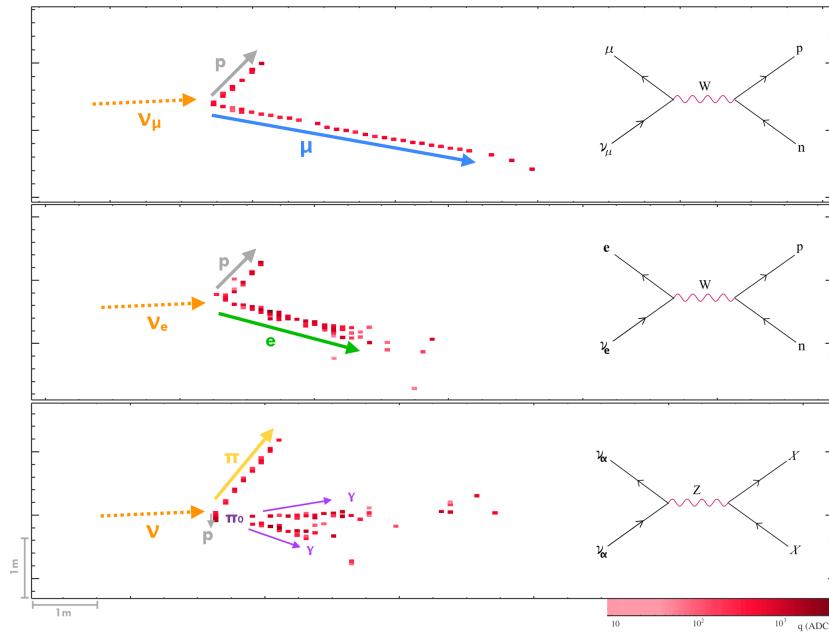
The suite of software we use is called NOvASoft, written primarily in C++ and Python, built on top of the Art data processing framework [84]. It is divided into a series of stages called modules that progressively add more information into the files in the form of new objects or “data products”, built using previous stages products as inputs. These products represent topological features in the events such as tracks, or contain the outputs of algorithms that estimate particle type or energy.

The data and MC files have the same underlying format and are passed individually through the same stages of the reconstruction. The simulation files have additional access to truth information from the generators, which can be used to assess the efficiency of algorithms. Neutrino events will look similar between the FD and ND due to both detectors having the same composition, so the same software chain can be applied to files from either detector as well.

We can divide the full reconstruction chain into three main parts: event reconstruction, energy estimation, and particle identification. We will now go through each part in detail.

## 4.1 Event Reconstruction

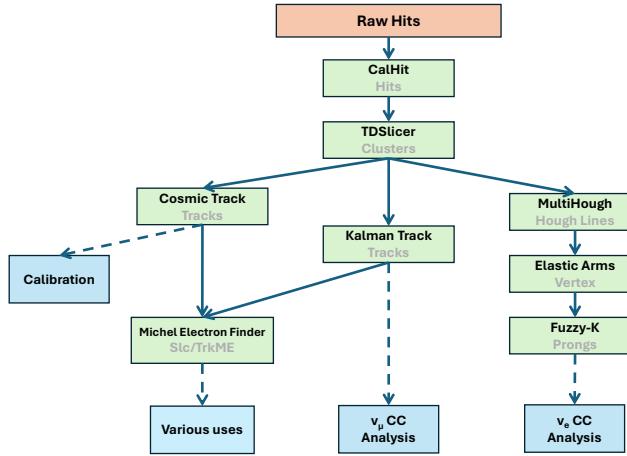
Event reconstruction builds the neutrino events out of individual hits in the detector, using algorithms designed to look for the topological features that different particles are known to make in the detector.



**Figure 4.1:** The typical topologies for  $\nu_\mu$  CC (top), and  $\nu_e$  CC (middle) events in our detectors, and an example NC (bottom) topology that appears  $\nu_e$ -like. The arrows indicate the particle type that produced the hits, which are shown in red. Each square represents a cell in the detector, and the color indicates the amount of charge, or energy deposited in that cell, with darker cells having more energy. Images from [85].

Figure 4.1 shows an example of what different neutrino interactions look like in our detectors. The first is the  $\nu_\mu$  CC interaction which is characterized by a long straight muon track. There may also be a shorter track close to the interaction vertex from hadronic activity, often a final-state proton. The  $\nu_e$  CC events will instead produce an electron as the charged lepton, which results in a shorter and broader electromagnetic shower. This event also contains a proton track near the interaction vertex. The neutral current event shown produces a  $\pi^0$  instead of a charged lepton, along with additional hadronic activity near the vertex. The  $\pi^0$  decays to two photons that shower, which can resemble an electron

shower, particularly if the photons happen to overlap. To distinguish these from  $\nu_e$  CC events, the displacement of the photon shower’s vertex from the primary interaction vertex can be identified.



**Figure 4.2:** Block diagram showing the order that reconstruction chain modules are run. The red block is the initial raw data input to the chain, the green blocks are the modules, with the names in black and data products produced in gray. The blue blocks show the eventual end usage of each part of the chain.

Figure 4.2 shows the reconstruction paths for the modules, and the end usage of the data products that were created. The green boxes represent individual modules that are run, each designed to find a specific feature in the data, and save it as an object into the file. Recall from Sec. 2.3.3 that data is organized in the files by run, subrun, and beam spill. Modules are run once per beam spill, and can access any data products associated with that spill that were created earlier on in the chain. We now discuss these modules in turn.

### 4.1.1 Hits

The first step in the reconstruction chain is to turn the collection of raw hits from our detectors into “calhits”. The coordinates of each raw hit are translated from DAQ logical channel number to the plane and cell number of the hit. The plane number gives the longitudinal position (Z-coordinate) in the detector, as well as which view the hit was in

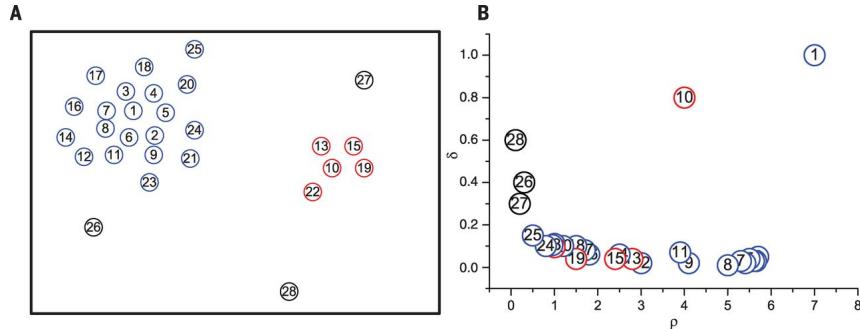
(X or Y), and the cell number gives the transverse position along that view. No distinction is made yet between hits from physics activity and those from electronics noise.

#### 4.1.2 Clustering

Next, calhits are clustered together in space and time. These clusters are referred to as slices and represent individual physics event candidates. NOvA uses an algorithm called TDSlicer, which first clusters calhits in each view separately, before merging the sets of 3D ( $x, z, t$ , and  $y, z, t$ ) clusters from each view into one set in 4D ( $x, y, z, t$ ). It has three steps; first, in each view, centroids are identified by finding the hits with maximum density and isolation according to the work of Rodriguez and Liao [86]. Density is defined as

$$\rho_i = \sum_j \exp(-d_{ij}^2/\tau^2) , \quad (4.1)$$

where  $d_{ij}$  is the Euclidean distance between hits in the same view, and  $\tau$  is a configurable scale parameter, set to 12(60) ns in the ND(FD) to reflect the timing resolution for hits in each detector [87]. Isolation is defined as the Euclidean distance to the closest hit in the same view with a higher density. This process is illustrated in Figure 4.3.



**Figure 4.3:** Illustration of the centroid-finding algorithm utilized by TDSlicer, based on [86]. The right plot shows the density ( $\rho$ ) and isolation ( $\delta$ ) of points for the clustering problem on the left. The points with the highest density in each group (1, 10) will also have high isolation and are chosen as centroids, using detector specific cutoffs for  $\rho$  and  $\delta$ . Points with high density but low isolation will form the bulk of each cluster. Points with low density and high isolation will be treated as noise. Image from [86].

Second, Prim's algorithm [88] is used to build up hits into a cluster, starting from the centroids. This is done sequentially, adding the next hit outside the cluster that is closest to some hit in the cluster, up to a predefined distance in each detector. Finally, the set of 3D slices in each view are merged by comparing the average  $(z, t)$  coordinates between slices. Each slice is compared to every other slice in the opposite view, and the slice with the closest match in  $(z, t)$  (up to a specified distance) is paired. These two slices are merged and removed from the list, and the process continues until all slices have been considered. If a matching slice cannot be found within the defined range for each detector, the slice is considered noise.

Once a slice is made, it is treated as a physics event. This object is the basis of most of our reconstruction infrastructure. Since slices are meant to represent individual neutrino events, we often refer to slice and event interchangeably <sup>1</sup>. Features of each event are picked out using further modules and then used to identify the event.

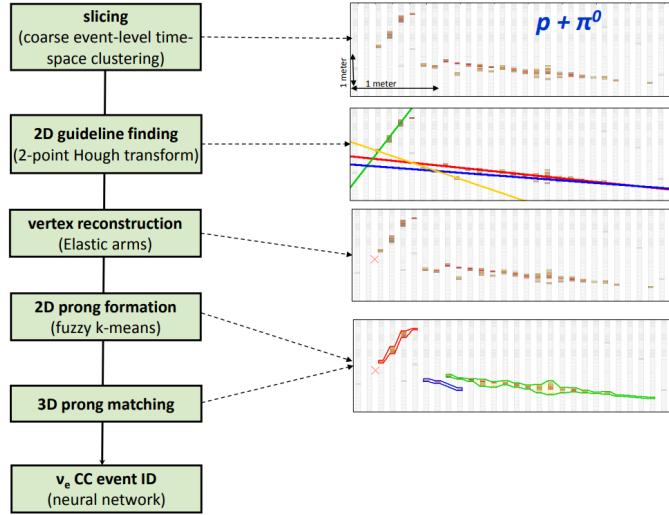
#### 4.1.3 $\nu_e$ reconstruction

The goal of the  $\nu_e$  reconstruction chain is to identify the electromagnetic shower produced by the electron in the  $\nu_e$  CC interaction. This is done through a series of modules, which first identify lines of interest, then use them to find the vertex of the interaction, and finally form prongs that extend from that vertex. This process is shown in Fig. 4.4.

The first step in the  $\nu_e$  chain identifies lines of interest in the clusters produced by TDSlicer by using a modified Hough transform [90]. This pattern-finding algorithm considers pairs of hits in each view separately, and calculates the line through them in polar coordinates  $\rho$ , the perpendicular distance from the origin to the line, and  $\theta$ , the angle between  $\rho$  and the axis of the view in question. The coordinates  $\rho$  and  $\theta$  from each line

---

<sup>1</sup>This is easily confused with an Art event, which is how the Art framework refers to one beam-spill worth of data from the detectors. I will specify “Art event” if using it in this context.



**Figure 4.4:** Flowchart of the  $\nu_e$  reconstruction chain applied to an NC event. Image from [89].

are then used to fill a 2D array, called a Hough Map, with a Gaussian vote given by

$$vote = e^{-\frac{(\rho-\rho_0)^2}{2\sigma_\rho^2}} e^{-\frac{(\theta-\theta_0)^2}{2\sigma_\theta^2}} \quad (4.2)$$

where  $\sigma_\rho = 3/\sqrt{12}$  and  $\sigma_\theta = 3/d\sqrt{6}$  ( $d$  being the distance between the two hits in the detector) [89]. After filling the Hough map, it is then smoothed by averaging votes in the map using a Gaussian smoothing weight. Prominent features in the event will lead to lots of votes in a particular region of the Hough map. Peaks above a threshold in the map are then identified, and the coordinates of those peaks used to define Hough lines that characterize the event.

This process is then iterated to search for smaller lines that could be missed by the presence of dominant lines in the map. Hits associated to the most dominant Hough line are removed, and the algorithm is re-run to identify new peaks in the map, which are now more likely to be associated with legitimate shorter physics tracks instead of noise [89]. This “Multi-Hough” process is repeated, keeping the dominant lines at each iteration, until

a maximum number of lines is reached or no peaks above threshold are identified in the map.

The next step attempts to find the neutrino interaction vertex, using the Hough lines as input. A modified version of the Elastic Arms method is used [91]. This algorithm assumes an event containing 1 or more tracks emanating from a common vertex. Taking the Hough lines and their intersection points as seeds, vertex points are identified, and 2D vectors, or “arms” are constructed pointing outwards from the vertex. The number of arms is taken as the larger of the number of Hough lines in either view. The direction of the arms, along with the vertex location are then adjusted to minimize an energy cost function [89]. This function rewards arms that pass through hits, and penalizes for missing hits. The version used in NOvA has an additional term which penalizes arms whose first hit occurs beyond one radiation length. This is necessary since, unlike in the traditional algorithm, the vertex is not known *a priori*, and also allows for the displaced photon showers from  $\pi^0$  decay common in NC backgrounds. The vertex which minimizes the energy cost function is then used.

Once the vertex is established, the next step is the assignment of hits in the event to “prongs”. Prongs are collections of hits meant to represent individual particle tracks or showers emerging from the vertex. Prongs are formed using a fuzzy k-means algorithm. “Fuzzy” refers to the fact that hits can be assigned to multiple prongs, and the total number of prongs is unknown at the start. In each view, prong centers are identified by observing the angular distribution of energy depositions around the vertex, and searching for peaks. Hits are then assigned a degree of membership to the prongs based on their proximity to the centers, and hits with a low degree of membership to any prong are treated as noise.

The final step is to merge the prongs together from each view. This is done by matching prongs which share a similar energy profile. A temporary track is made out of every prong in each view, and the cumulative energy distribution along the track is computed. This energy profile for each prong is then compared to those in the opposite view using a Kuiper metric [89], which takes the sum of the absolute values of the largest positive and negative

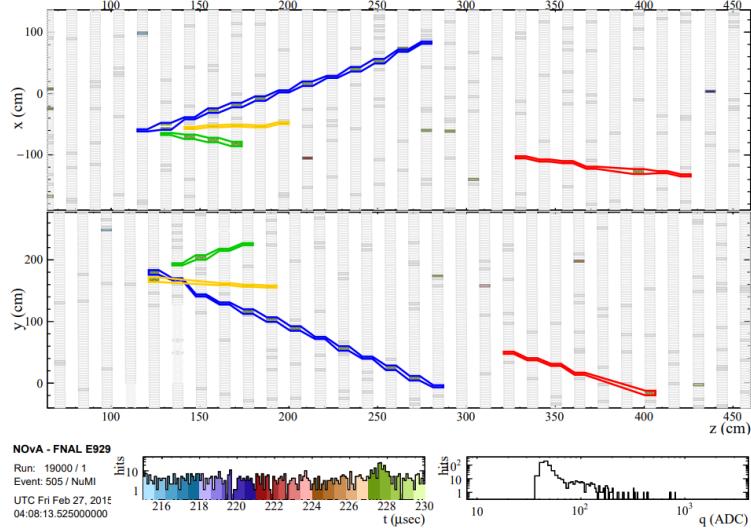
vertical distances between the two distributions. The pair of prongs with the minimum value of this metric are matched, and removed from the list of available prongs. The process is repeated until all prongs that can be matched are matched.

#### 4.1.4 $\nu_\mu$ Reconstruction

The  $\nu_\mu$  reconstruction chain is more straightforward, owing to the highly identifiable track of hits that muons make in the detector. Instead of showering, the muons lose energy via ionization, tracing a straight and narrow path through the detector with occasional small changes in trajectory from multiple scattering. The module we use to identify these and other tracks is called Kalman Track [92], which makes use of a Kalman Filter [93] to identify lines with possible deviations.

The algorithm starts by looking in the downstream end of the detector for pairs of hits in the same view separated by no more than 4 cells, to be used as track seeds. The tracking begins downstream since the particles created in the interaction are expected to be more separated there. It uses the track seeds to calculate an initial slope. The algorithm then scans back and forth through the detector, iteratively adding hits to build up the tracks. The location of adjacent hits is predicted using the estimated position and direction of the track assuming a linear fit to the current track, with some weighting to account for scattering and measurement uncertainty. Hits that are within 8 units of  $\chi^2$  from this track prediction are added to the track [92]. After adding each hit the estimated trajectory is updated, and the process continues until no valid hits can be found.

The 2D tracks from each view are merged together to form 3D tracks. As a first step, tracks in a single view are compared to look for broken tracks that can be joined together. This can occur if there were gaps in the hits, or if a hard scattering occurred that produced a large angular deviation, leading to two independent tracks being made for one particle. The final list of tracks in each view are then merged together based on their overlap in the Z direction, resulting in 3D tracks. An example of the tracks reconstructed for a simulated  $\nu_\mu$  event can be seen in Fig. 4.5.



**Figure 4.5:** A zoomed-in FD event display showing the Kalman tracks reconstructed for a simulated  $\nu_\mu$  event. From [92].

#### 4.1.5 Cosmic Tracks

CosmicTrack is the module used to identify the downward-going tracks from cosmic activity in both detectors, assuming a simple topology of a single straight-line track in the detector, typically at a steep vertical angle. It is designed as a simple, fast reconstruction method for these single-particle events which are mostly used for calibration (Sec. 3.2). It uses a sliding window-tracking algorithm that is performed separately in each view before merging views [94].

For each reconstructed slice, a user-defined range of planes (the “window”) at one end of the slice is selected. The hits within that window are used to fit a straight line, and any hits in the slice consistent with that line are added to the track. Then, the window is moved one plane forward in the slice, and any new hits consistent with the line are added. A new fit is then performed to update the line’s trajectory with the new hits. This process is repeated until all planes in the slice have been considered. The tracks in each view for each slice are then merged together making a final 3D track. These tracks are primarily used for calibration, as discussed in the previous chapter.

#### 4.1.6 Michel electrons

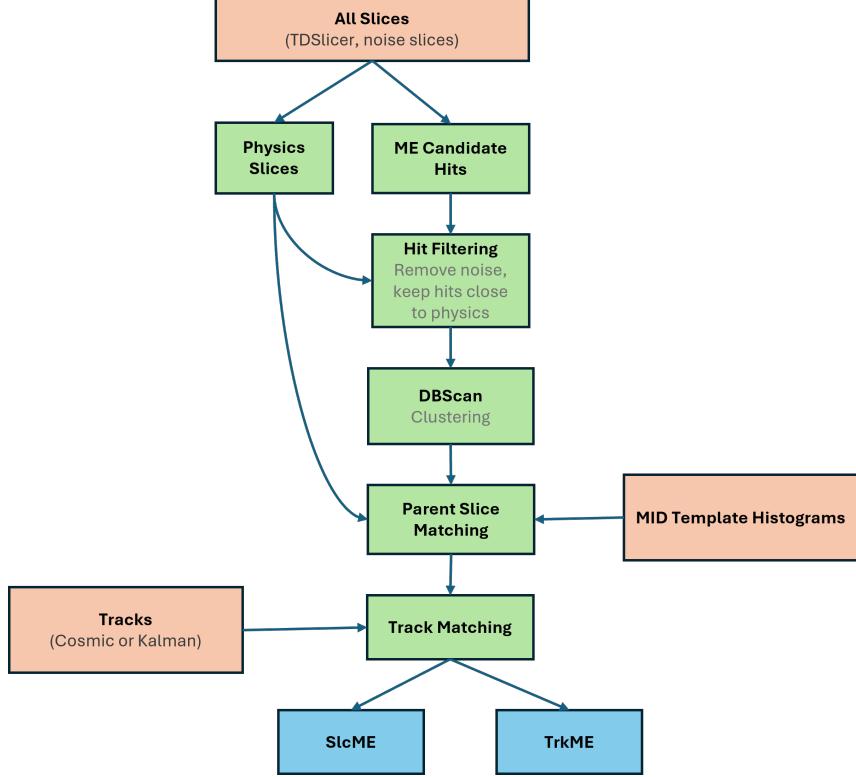
A Michel electron is the electron that results from a muon decay ( $\mu^- \rightarrow e^- + \nu_\mu + \bar{\nu}_e$ ). It shows up in our detector as a small time-delayed signal near the ends of a muon track. Michel electrons are useful for tagging muons that stop in the detector, and they provide a source of events with well-understood energy and decay-time distributions. They have several uses for NOvA, including tagging stopped cosmic muons for calibration (discussed in Sec. 3.2), understanding electron neutrino beam backgrounds through our decomposition process (Sec. 5.4.1), or as a standard candle to set our energy calibration uncertainty (Sec. 5.5).

Because of the possibly long time delay between the muon stopping and electron appearing, hits from the Michel electron are unlikely to be clustered with the hits from their parent muon. Additionally, Michaels have an energy distribution that cuts off at 53 MeV (half the muon rest mass), so they do not deposit a lot of energy or make a lot of hits in the detector. Our goal is to both find and cluster the small hits from the Michel, and then match it with a parent slice (the muon).

Michel Electron Finder (MEFinder) is the module we use to do this [95]. It is run once per Art event, and does the Michel clustering and matching across all hits and slices in that spill at once. Figure 4.6 shows an overview of the steps in this process.

We start by selecting the appropriate hits to cluster. Any detector hits that did not end up clustered by TDSlicer are added to a “noise slice” for each Art event. We form a list of candidate Michel electron hits by considering all hits from the noise slice, as well as all hits from slices from TDSlicer with a low total number of hits ( $<=12$ ). We refer to the rest of the slices from TDSlicer with  $> 12$  hits as “physics slices”.

We then filter the sample of candidate Michel hits. We include only hits with ADC values  $>= 50$  to remove electronics noise. Since we expect the Michel to show up as a delayed signal near physics slices, we only keep hits within 10,000 ns of the mean time, and 40 cm of a physics slice. Additionally we remove any hits if another hit occurred



**Figure 4.6:** Block diagram showing the sequence of steps in the MEFinder algorithm. The arrows indicate the order of steps performed, and where the outputs from each step are used. Red blocks indicate inputs from outside the algorithm, green blocks are the steps performed by the algorithm, and the two classes of output objects, SlcME and TrkME, are shown in blue.

in the same cell within 1200 ns. This is to account for a data/MC discrepancy in APD deadtime. The APD deadtime had been mismodeled in the simulation, leading to a lower average deadtime, and some simulated hits occurring too soon after one another<sup>2</sup>. This cut removed such hits as a way to give a better Data/MC agreement without needing to alter the low-level simulation and require remaking all of the files.

Once we have the list of candidate Michel hits, MEFinder makes clusters of hits using a simplified version of the density-based clustering algorithm DBScan [96], taking the plane, cell, and time differences of hits as input. The algorithm calculates the distance between

---

<sup>2</sup>This is discussed further, along with a solution for future analyses, in Sec. 6.4.

hits using the distance function

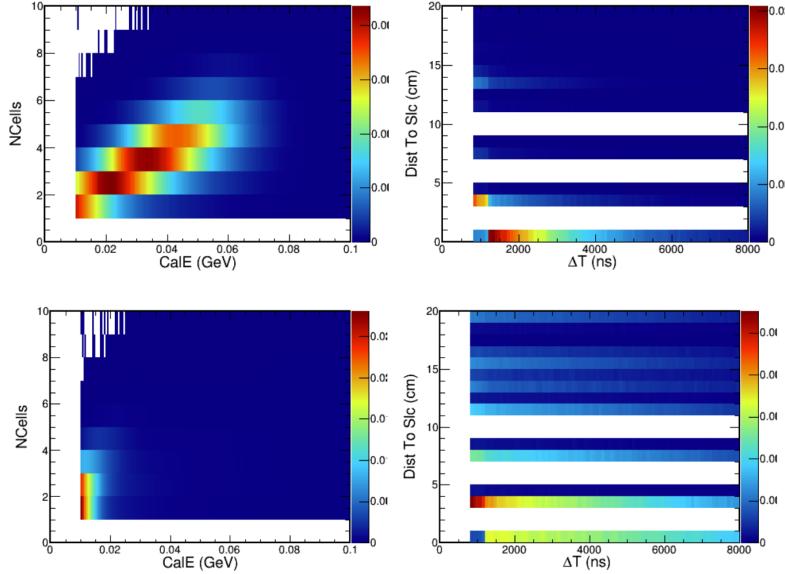
$$d = \sqrt{\left(\frac{\Delta P}{1.85}\right)^2 + \left(\frac{\Delta C}{1.40}\right)^2 + \left(\frac{\Delta T}{40\text{ns}}\right)^2} \quad (4.3)$$

where  $\Delta P$  is the difference between plane number of the hits,  $\Delta T$  is the time separation, and  $\Delta C$  is the difference in cell number when the two hits are in the same view, otherwise  $\Delta C = 0$ . Starting with individual hits, clusters are built up by adding hits that are within 5 distance units of any other hit in the current cluster. Clusters from the DBScan output that contain no more than 12 hits, and are within 10,000 ns and 20 cm of a physics slice will become a Michel cluster.

The next step is to determine which physics slice to assign as the parent slice for each Michel cluster. Each Michel is matched to only one physics slice. For most clusters, there will be only one physics slice within the allowed range, and that will be matched. For Michaels with multiple possible parents we need a way to determine the best match. This is done using our Michel electron Identifier (MID). MID is a log-likelihood particle identifier that we use to select true Michel electrons, and match parent slices accurately. It is built on four reconstructed input variables that look significantly different between true-ME and non-ME clusters. These are the reconstructed calorimetric energy (CalE), the number of hits (NCells), the time difference between the parent slice and the Michel ( $\Delta T$ ), and the distance between the parent slice and Michel (DistToSlc). Normalized 2D template histograms of these variables are made from simulation files, using truth information to separate the true-ME and non-ME components. These are shown in Fig. 4.7. The parent slice assigned to each Michel cluster is chosen to be the physics slice that would give the highest value on the true-ME DistToSlc vs.  $\Delta T$  template.

Once the parent slice is matched, we can calculate the MID, which represents the likelihood that the cluster is from a true Michel electron. This is done using the formula

$$MID = \log(L_{true}) - \log(L_{non}) \quad (4.4)$$



**Figure 4.7:** Normalized 2D template histograms of the Michel variables used in the MID calculation. The top two plots are for the true-ME case, and the bottom two are for the non-ME case. The empty bands in the DistToSlc plot come from the distance metric used, which takes the closest cell-to-cell distance between Michel and parent slice hits in the same view. The cell-to-cell distances are roughly standard, so only certain values are allowed, hence the bands. True Michels have more energy and hits on average, and occur in the same cell as the parent slice ( $\text{DistToSlc} = 0$ ) more often. From [95].

where  $L_{true}$  and  $L_{non}$  are the likelihoods for the true/non-ME hypotheses. These are calculated using the probabilities pulled from the MID templates for a given Michel cluster's input variables ( $\text{CalE}$ ,  $\text{NCells}$ ,  $\text{DistToSlc}$ , and  $\Delta T$ ) according to

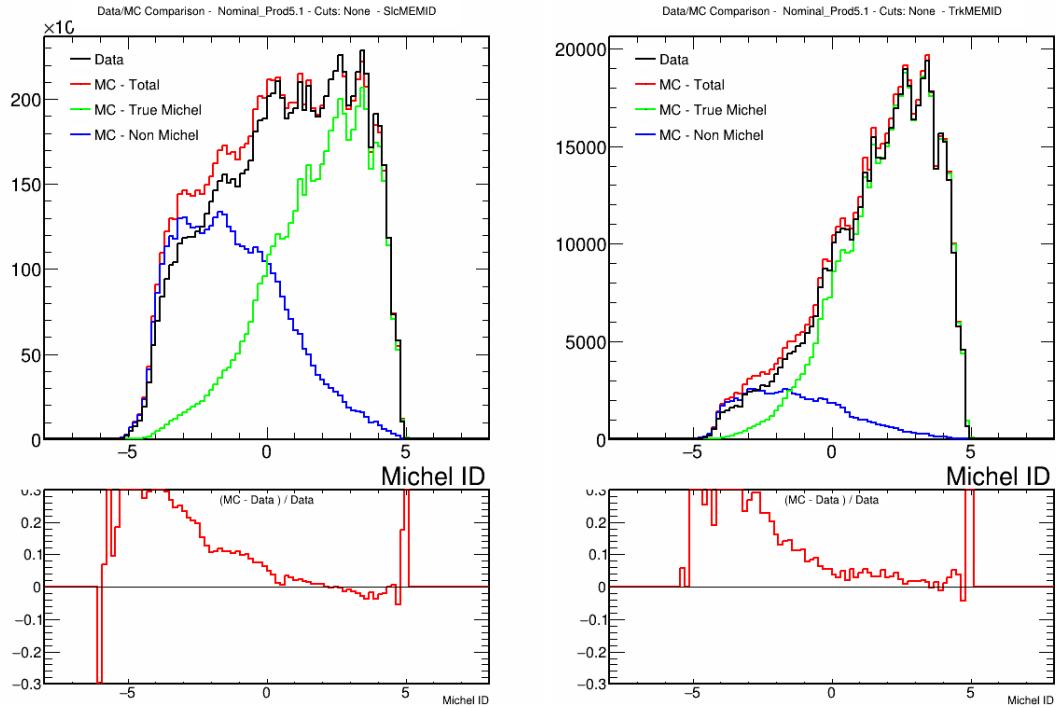
$$L_{true} = P_{true}(\text{CalE}, \text{NCells}) * P_{true}(\text{DistToSlc}, \Delta T) \quad (4.5)$$

$$L_{non} = P_{non}(\text{CalE}, \text{NCells}) * P_{non}(\text{DistToSlc}, \Delta T) . \quad (4.6)$$

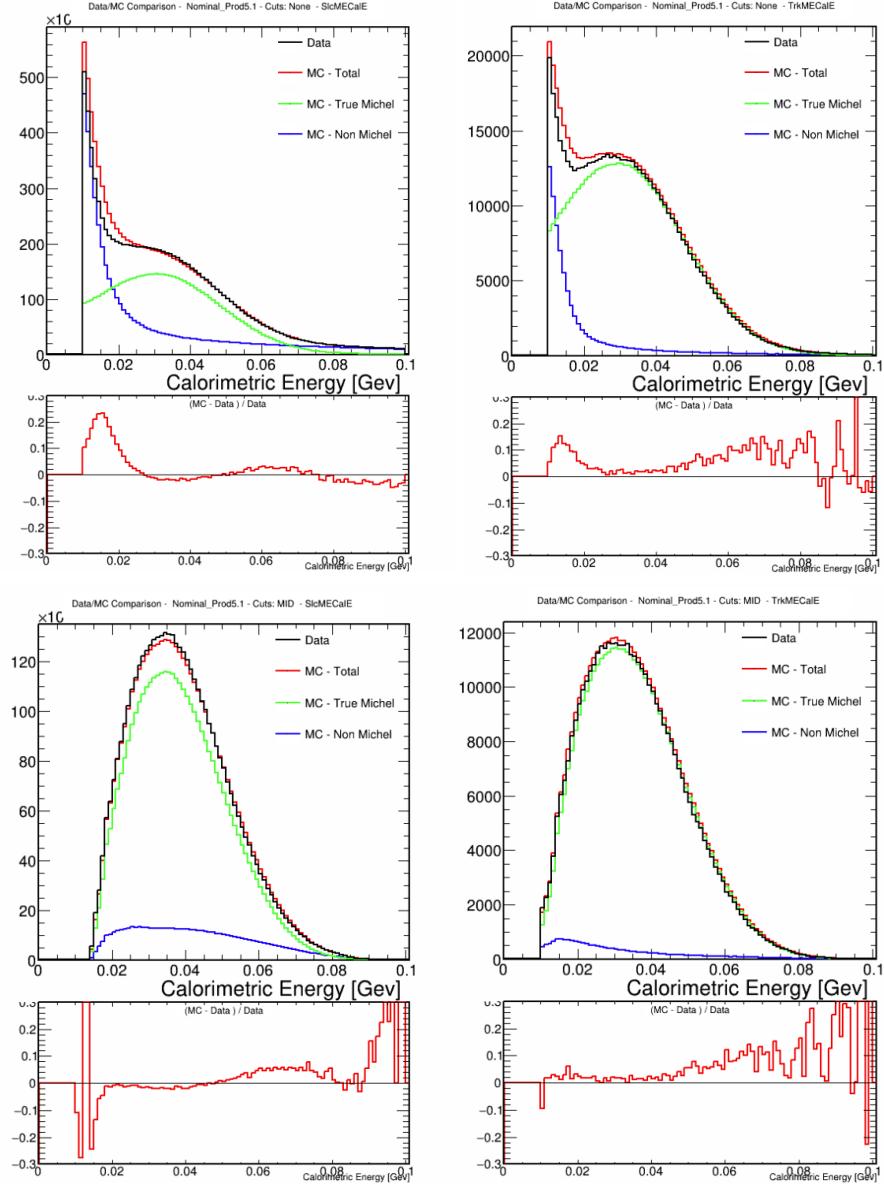
MEFinder outputs two classes of Michel electron objects, SlcME and TrkME. A Michel cluster is considered a TrkME if it is within 15 cm of some Kalman (Cosmic) 3D track endpoint with track length  $\geq 100$  cm. This Michel cluster is then matched to the track and its associated slice. All other Michel clusters are labeled as SlcME. TrkMEs represent a purer sample of Michel electrons given the extra requirement of the presence of a muon

track. As such they are useful for precision checks in the calibration and energy scale calculations.

Data/MC comparisons for the distribution of MIDs of SlcME and TrkME in the ND are shown in Fig. 4.8. The effect of cutting on MID can be seen in Fig. 4.9, which compares the energy distributions of the Michels.



**Figure 4.8:** Data/MC comparison of the MID values of reconstructed SlcME (left) and TrkME (right) in the ND. The top plots show the individual data and MC distributions, while the bottom shows the ratio  $\frac{MC - Data}{Data}$ .



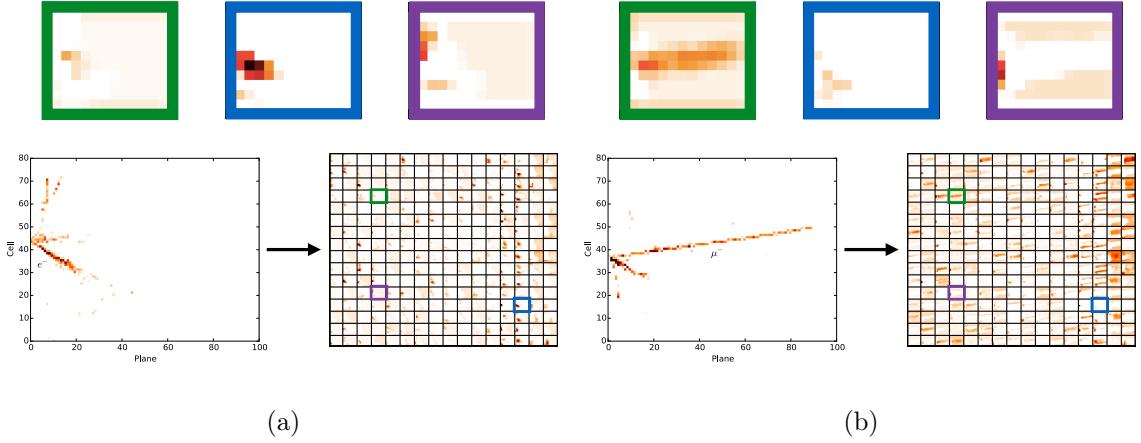
**Figure 4.9:** Data/MC comparison of calorimetric energy for reconstructed SlcME (left) and TrkME (right) in the ND for FHC events. The top plots are the distributions without any cuts applied. The bottom distributions have a cut on  $MID > 1$  applied to SlcME and  $MID > 0$  applied to TrkME. Beneath each plot is the ratio distribution  $\frac{MC - Data}{Data}$ .

## 4.2 Particle Identification

NOvA makes use of various machine learning techniques to identify and separate  $\nu_e$  CC,  $\nu_\mu$  CC, NC, and cosmic events, using the reconstructed objects as inputs. The algorithms will output scores, which represent the likelihood that an event is of a particular neutrino type, or that individual tracks/prongs are from a particular particle type. These scores are later used to develop cuts to select different samples (see Sec. 5.2). We will describe the algorithms that are most relevant for this analysis.

### 4.2.1 Event Classification with CVN

The Convolutional Visual Network (CVN) is a deep neural network that is used as our primary event classifier [97]. It is based on Convolutional Neural Networks (CNN), which are a class of algorithm typically used for image recognition tasks. In the case of CVN, the “image” is a pixel map of the slices, where each pixel represents a cell, and the intensity represents the amount of energy deposited in that cell. Pixel maps are generated for each view in the slice, and input to the algorithm separately, before being combined at a later stage. Training CVN on the slices avoids reconstruction failures such as broken tracks or misplaced vertices, and allows the neural network to classify events based on more complex or abstract topological features that may not be expected by humans. The features identified by the algorithm can be visualized in a “feature map”. Figure 4.10 shows examples of pixel maps for  $\nu_\mu$  and  $\nu_e$  CC events, as well as the feature maps for each. CVN is trained over millions of simulated beam events, and hundreds of thousands of cosmic events [98]. The output of CVN is a set of scores from 0 to 1, which represent the likelihood that an event is  $\nu_\mu$  CC,  $\nu_e$  CC, NC, or cosmic. Each event’s individual scores for each category will add up to 1, giving the probability the event is of that type. Figure 4.11 shows the distribution of  $\nu_\mu$  and  $\nu_e$  scores by true event type.

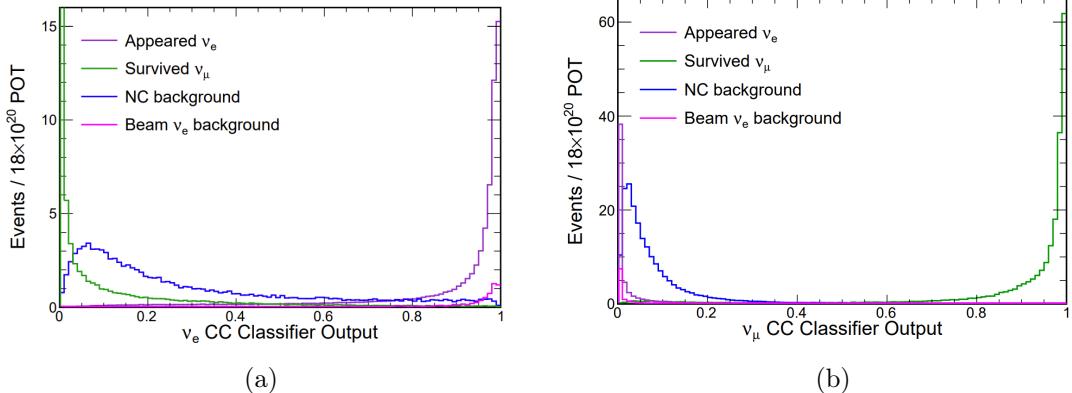


**Figure 4.10:** Pixel map of the Y-view of  $\nu_e$  CC (left), and  $\nu_\mu$  CC (right) events along with the feature maps extracted for each. Individual features are highlighted corresponding to muon-like (green), electron-like (blue), and hadronic (purple) qualities in each event. Images from [99].

#### 4.2.2 Muon Identification using ReMID

To better identify the muon tracks from  $\nu_\mu$  CC events we use the Reconstructed Muon Identifier (ReMID) [100]. ReMID is a Boosted Decision Tree (BDT) which takes the Kalman tracks from each event as input and scores them based on how muon-like they are. ReMID is primarily meant to distinguish muons from their dominant background, charged pions from NC events, and makes use of variables that look significantly different between the two particles.

The four variables input to the BDT are the  $\frac{dE}{dx}$  log-likelihood, the scattering log-likelihood, the track length, and the fraction of planes that overlap with hadronic activity. The energy deposited along the track, or  $\frac{dE}{dx}$ , can be different between muons and pions, because pions will lose energy from hadronic scattering in addition to ionization. Energy profiles for muon and pion hypotheses are made, and a log likelihood is computed for each track comparing its own  $\frac{dE}{dx}$  to these templates. The hadronic scattering will also lead to differences in the angular deviations of pions as a function of track length. The scattering log-likelihood uses an analogous process to determine how much the track differs from a



**Figure 4.11:** Distribution of CVN output scores for  $\nu_e$  CC (left) and  $\nu_\mu$  CC (right) for simulated FD events, split by true interaction type. Both are able to effectively separate signal and background, with the  $\nu_e$  classifier having a small irreducible background of  $\nu_e$  CC events that did not oscillate from  $\nu_\mu$ , but were intrinsic to the beam. The  $\nu_\mu$  performance is boosted by the ease with which we can identify muon tracks in our detector. Images from [97]

straight line. The track length is actually similar between muons and pions, except at higher energies where muons produce longer tracks. This variable is primarily meant to distinguish muon tracks from hadronic showers that produce many particles with short tracks. Finally, the fraction of planes in the track that overlap with hadronic activity is expected to be very low for muons except near the interaction vertex. The opposite is true for pions, whose interactions can create additional hadrons along the track.

The highest scoring ReMID track in each event is later used in the  $\nu_\mu$  CC event selection process (section 5.2).

### 4.2.3 Prong CVN

Prong CVN is an additional application of the CVN architecture to identify the particle type of individual prongs within an event. It differs from event CVN by having four total pixel map inputs: two generated from the hits in each view of the entire slice (similar to event CVN), and two smaller maps generated from the hits in each view of the individual prong. The output is again a set of scores ranging from 0 to 1, representing the probability

that an individual prong was the result of an electron, proton, photon, muon, or pion. For this analysis these scores are only used for identifying the hadronic vs. electron components of the slice during  $\nu_e$  energy estimation (Sec. 4.3.1).

#### 4.2.4 Cosmic Rejection BDTs

As an additional step to reject cosmic backgrounds at the FD, we use Boosted Decision Trees to score events based on how likely they are to not be cosmic in origin. There are separate BDTs trained for the  $\nu_\mu$  and  $\nu_e$  samples, but using similar kinematic input variables. In both cases the BDTs are trained using simulated beam events as the signal, and cosmic data for the background.

The  $\nu_\mu$  BDT uses the following variables: the track length, the cosines of the track angle with respect to the beam and Y axis, the start and end positions of the track, the distance from the track start/end to any wall of the detector, the number of hits in the track vs. slice, and the  $p_t/p$  or transverse momentum vs. total momentum for the track.

The  $\nu_e$  BDTs are more complex. As will be described in Sec. 5.2, we split our  $\nu_e$  events into core and peripheral sub-samples, which represent fully contained and partially contained  $\nu_e$  events, respectively. Due to the different geometries of these events, we use separate cosmic rejection BDTs for each. The BDT for the core  $\nu_e$  sample is trained using a similar set of variables to the  $\nu_\mu$  sample, but using prongs instead of tracks. The BDT for the peripheral  $\nu_e$  sample is trained using a smaller set of variables, using only  $p_t/p$ , the vertex  $x$ ,  $y$ , and  $z$  positions, and the distance of the prong start/end to any wall of the detector other than the top.

The BDTs perform very well, with >90% cosmic backgrounds rejected while retaining at least 90% signal efficiency in all cases [101, 102].

#### 4.2.5 Low-Energy $\nu_e$ BDT

New for this analysis is an additional sample of reclaimed low-energy  $\nu_e$  events that fail our main selection cuts. Their energy range, from 0.5 - 1.5 GeV is one where we typically

have low selection efficiency, and higher backgrounds, so a simple retuning of the selection cuts was not possible to reclaim these events [103]. Instead, a BDT was developed to select these events, being trained on both hit-level reconstruction variables, prong topology and energy variables, and CVN scores [103]. The training sample used simulated events that fail our main  $\nu_e$  selection, outlined in Sec. 5.2. For training, all  $\nu_e$  events were considered signal, with no distinction made between beam  $\nu_e$  and oscillated  $\nu_\mu \rightarrow \nu_e$  events. All other categories of events were considered background.

## 4.3 Energy Estimation

To constrain the oscillation parameters requires observing neutrino events as a function of energy at a fixed baseline. This in turn requires an accurate measurement of the energy of each neutrino event. The first step in energy estimation, the calibration, was already discussed in the previous chapter. Once the neutrino events have been reconstructed, we can employ specific energy estimation methods for  $\nu_e$  and  $\nu_\mu$  events that achieve better performance than calorimetry alone.

### 4.3.1 $\nu_e$ energy estimation

The energy of  $\nu_e$  events is divided into components from the hadronic activity, and the electromagnetic shower produced by the electron. The  $\nu_e$  energy is reconstructed from

$$E_\nu = p_0(p_1 E_{EM} + p_2 E_{had} + p_3 E_{EM}^2 + p_4 E_{had}^2) \quad (4.7)$$

where  $E_{EM}$  and  $E_{had}$  are the energies of the electromagnetic shower, and hadronic components respectively. These are obtained by summing the visible energy of EM and hadronic prongs, which are selected using the Prong CVN. The parameters  $p_i$  are found by fitting the polynomial to a weighted-average true energy distribution. The weights serve to flatten the true energy distribution to avoid biasing towards the peak energies [104]. The fit is

performed separately for FHC and RHC events. The overall energy resolution for  $\nu_e$  events is 10.8% for FHC, and 8.5% for RHC [105].

### 4.3.2 $\nu_\mu$ energy estimation

We can calculate muon energy based on the distance a track traverses in the detector. The uncertainty on the muon track length is on the order of one plane length, and lower than the uncertainty on our energy scale, allowing us to achieve better energy resolution than from calorimetry alone. The hadronic energy is then simply estimated by summing the visible energy of hits not associated with the muon track. The  $\nu_\mu$  CC energy can then be estimated by the equation

$$E_\nu = E_\mu + E_{had} \quad (4.8)$$

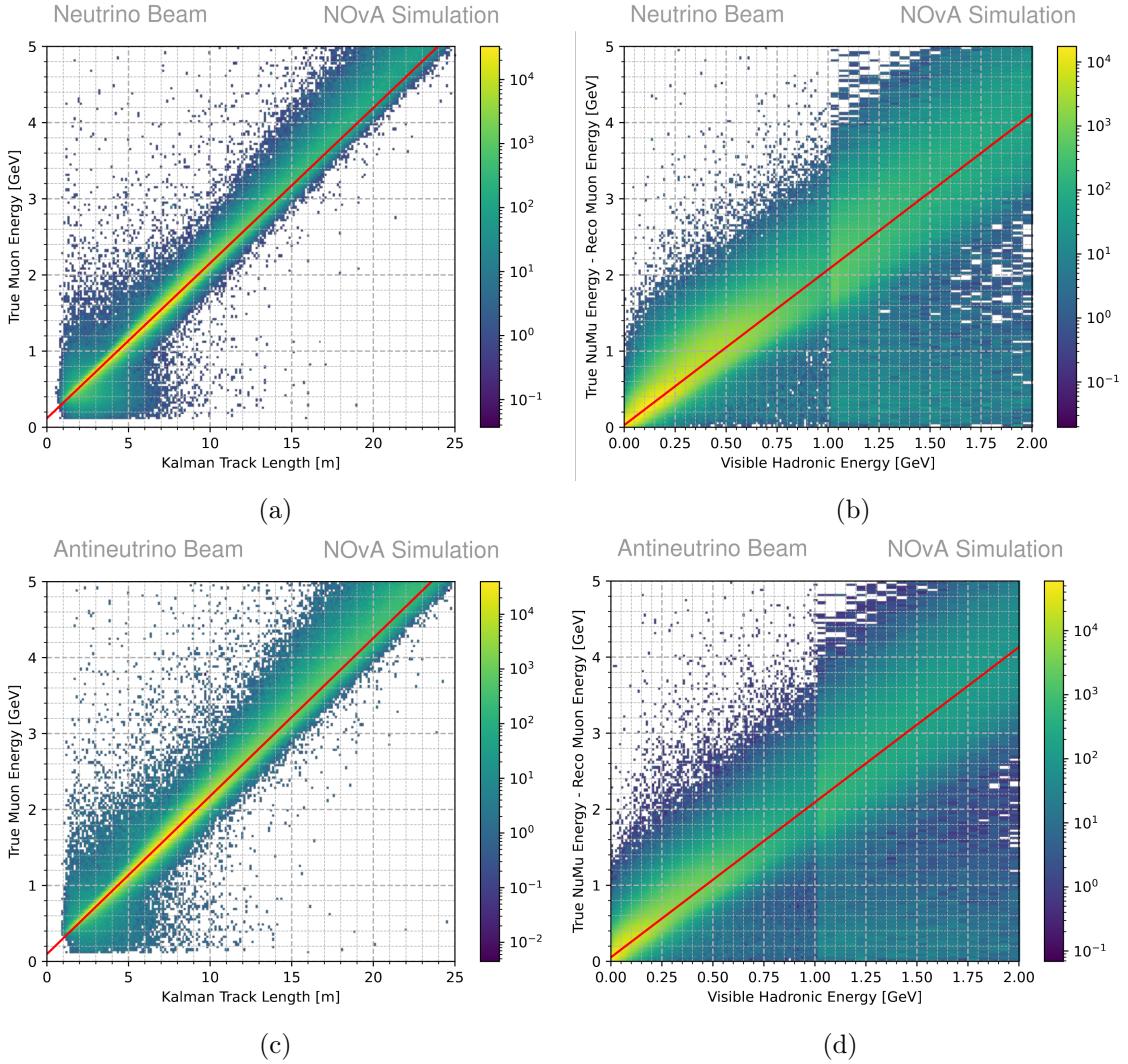
where  $E_\mu$  is the energy of the muon, and  $E_{had}$  is the energy of the hadronic component of the interaction.

Both muon and hadronic energies are corrected using fits to simulated events, shown in Fig. 4.12. 2D Histograms are made of the reconstructed quantity (Kalman track length, or visible hadronic energy) vs. true energy. In each bin of the reconstructed variable, the mode of the distribution of true energy is found, and the tails under 10% are removed. This modified true energy distribution in each bin is then fit to a Gaussian.

The means of the true-energy Gaussian fit in each reconstructed bin are then used as data points to perform a linear fit<sup>3</sup>. The points are additionally weighted by the size of the peak of the Gaussian distributions, to avoid losing information about the relative frequency of events in each bin. The result of the fit is shown as the red line in Fig. 4.12. The resulting energy resolutions for  $E_\nu$ ,  $E_\mu$ , and  $E_{had}$  are summarized in Table 4.1.

---

<sup>3</sup>In previous analyses, a piece-wise linear spline fit was used but it was found that this resulted in unwanted sharp features in the hadronic energy distribution due to fitting a piece-wise linear function to a continuous distribution. A simple linear fit gave very similar performance without these unwanted features [106].



**Figure 4.12:** Linear fits to the true energy (shown as a red line) for FHC (top) and RHC (bottom) events for the muon (left) and hadronic (right) components of the event. The higher energy resolution of the muon component is evident when comparing the spread of the distributions [107].

$\nu_\mu$ Energy Resolution			
Detector	Beam	Bias	Resolution
FD	FHC	-2.58%	9.30%
FD	RHC	-1.68%	8.06%
ND	FHC	-3.06%	12.18%
ND	RHC	-2.85%	9.98%
$\mu$ Energy Resolution			
Detector	Beam	Bias	Resolution
FD	FHC	0.02%	4.17%
FD	RHC	-0.07%	4.04%
ND	FHC	-0.00%	4.80%
ND	RHC	-0.14%	4.44%
Hadronic Energy Resolution			
Detector	Beam	Bias	Resolution
FD	FHC	-4.23%	31.11%
FD	RHC	-3.51%	34.10%
ND	FHC	-4.81%	38.59%
ND	RHC	-11.31%	42.42%

**Table 4.1:** Energy biases and resolutions for  $\nu_\mu$  events as well as individual  $\mu$  and hadronic components. These are calculated as the mean and standard deviation of the  $\frac{E_{reco} - E_{true}}{E_{true}}$  distribution [105].

# Chapter 5

## 3-Flavor Analysis Setup

### 5.1 Introduction

Our measurement of the oscillation parameters is extracted by fitting the energy spectrum of our simulated FD neutrino events to the observed data. We have so far described the process of reconstructing the events, estimating their energy, and assigning them particle ID scores to gauge their likelihood of being a particular flavor. All the basic pieces are now in place to carry out the analysis. We begin by developing cuts to select pure samples of  $\nu_e$  and  $\nu_\mu$  events, and binning them in ways that improve our sensitivity to the oscillation parameters when fitting. With these cuts and binnings we can construct predictions of the oscillated FD energy spectra, and then apply data-driven corrections using information from the ND. This is a two-step approach where the ND simulation is corrected to better match the data, in a process called decomposition, and then these corrections are propagated to the FD in a process called extrapolation. A wide range of systematic uncertainties affecting the predictions are also calculated. The final predictions and systematics are then used to test hypotheses of different oscillation parameters, using either a Bayesian or a Frequentist approach to inference. This chapter will cover each step in detail.

## 5.2 Event Selection

We apply a series of selection cuts to our data to get pure samples of  $\nu_e$  and  $\nu_\mu$  events for fitting. Similar categories of cuts are applied to each sample but many are optimized independently for  $\nu_e$  and  $\nu_\mu$  selection. The broad categories of cuts are:

- **Spill and Data Quality** - Ensures that both the beam and the detectors were running in good conditions. Additionally, it places timing cuts to only keep data in the beam spill window.
- **Containment** - Removes events whose particles are not fully contained in the detector. This ensures all of the neutrino energy was deposited in the detector, and can be fully reconstructed. Additionally, it helps to remove cosmic and rock muon background events.
- **Reco and Event Quality** - Applies cuts on the reconstructed objects and basic event variables such as the energy and number of hits. This checks that the events are well-reconstructed and fall within the desired energy range.
- **Particle Identification (PID)** - Uses the scores from our event CVN and BDTs to select the desired signal candidate events and remove cosmic and other backgrounds.

### 5.2.1 Spill and Data Quality

Occasionally, the beam or detectors undergo a failure mode, and the resulting data may not be of high enough quality to include in the analysis. These cuts remove poor quality data and are applied at either the subrun or individual spill level.

The data quality cuts applied at the subrun level compare various quantities averaged across the detector and subrun to label individual subruns as “good” or “bad”. There are several cuts a subrun must pass to be labeled as “good”. The average hit rate across the detectors, along with the rate of slices and 3D tracks produced must all be within nominal levels [108]. Timing quality cuts are also applied, checking that event timestamps

are ordered correctly, and hits are occurring in time with the beam. Finally, at least 4 diblocks in the FD, and all diblocks in the ND, must be reporting normal hit rates to allow for a minimum analyzable detector size. These cuts can catch failure modes associated with the hit rate, such as a change in gain from cooling issues, or time synchronization issues leading to broken tracks. Additionally, at the FD, if an individual APDs hit rate is producing too many hits ( $> 10^{3.5}$ ) or not enough hits ( $< 3$ ), it can be masked off and excluded from reconstruction [108].

Spill level cuts are meant to remove short periods of low-quality data that occur in an otherwise good run. There are two types of spill-level cuts. The first focuses on the beam quality. To pass, a beam spill must have a POT greater than  $2 \times 10^{12}$  and a total spill time of less than  $0.5 \times 10^9$  ns. In addition, the beam's position and width, and the focusing horn's current must all be within acceptable ranges [105]. Finally, the timing peak of the beam must be within the 217-229  $\mu$ s beam spill window of the 550 ns NuMI trigger. Only data within this window is used in the analysis. The second type of spill cut focuses on the detectors. These primarily catch failure modes where one or more DCMs are down or taking poor quality data. If a DCM in either detector is not reporting any hits, then that spill will be excluded. At the FD, there is an additional check on the timing of the DCMs, to ensure that they are not out-of-sync with each other [109].

The amount of data we keep for the analysis has increased over time. In the first data taking period, that occurred while the detector was still being built, the percent of POT retained after applying Good Runs and spill cuts was 73.3%. This increased to  $> 99\%$  in subsequent periods and remains high [110].

### 5.2.2 $\nu_e$ selection

The electron neutrino selection for the FD is outlined in the cut flow diagram in Fig. 5.1. The selection process results in several different samples of  $\nu_e$  events. The main sample is the “Core” sample, for which events must pass all quality, containment, and PID cuts. There are two additional samples, the “Peripheral” and “Low-energy” samples, which seek

to reclaim events into the analysis that fail the Core selection. Each has its own set of event quality and PID cuts tuned to the particular sample. The Peripheral sample is FD-only and focuses on events that were not fully contained in the detector, but were otherwise very  $\nu_e$ - like. The Low-energy sample is FHC-only<sup>1</sup>, and probes a new region of energy space for this analysis. It uses a separate BDT trained to identify low-energy  $\nu_e$  events that fail the Core selection. Additional details and motivation for this samples inclusion can be found in Chapter 6. The ND  $\nu_e$  selection follows a similar flow, but with its own cut tunings. The primary difference is the absence of a Peripheral sample and nearest slice cut.

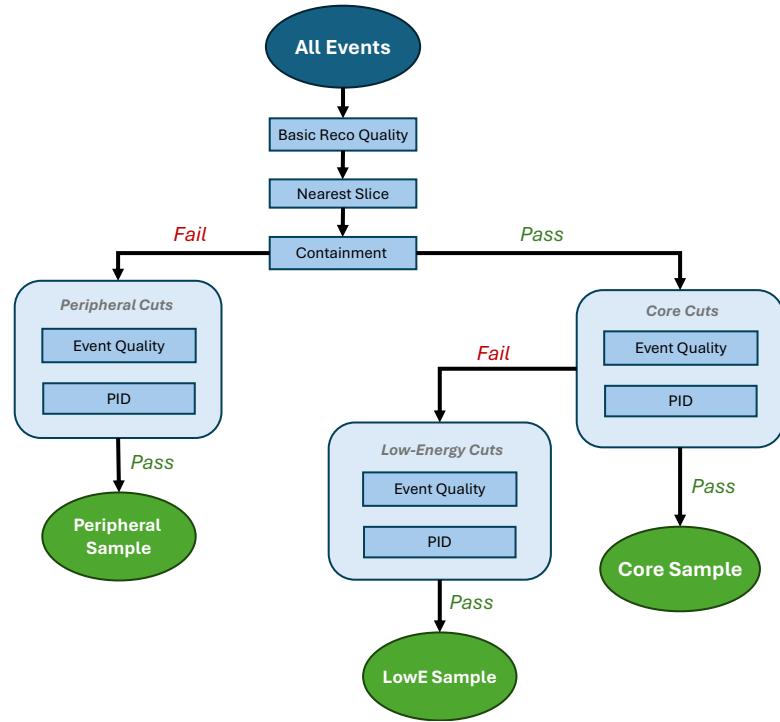
At the start of the cut flow are two cuts that all events need to pass. The basic reco quality cut checks that reconstruction was successfully run on the event. It requires a reconstructed vertex, and at least one prong in the event. It additionally requires there to be fewer than 8 hits per plane in the event, to remove a type of reconstruction failure called an FEB flasher. This occurs when a large amount of charge passes through one or more cells in a single FEB, saturating the electronics and causing all 32 channels on the FEB to “light up”. This is primarily observed in large cosmic background events.

The nearest slice cut requires the nearest slice in time to the candidate slice to be sufficiently far away, in either space or time, and sufficiently far from the top of the detector, with the exact ranges varying based on the CVN score of the candidate slice (lower scores requiring higher spacing). This cut reduces the Bremsstrahlung background from cosmic events coming from a reconstruction failure where a cosmic event is sliced into two separate events, with one of them appearing to be electron-like [111].

The containment cut defines the split between core and peripheral samples in the FD. To pass, an event must have all hits at least 63 cm from the top, 18 cm from the front or back, and 12 cm from the east, west, or bottom sides of the detector. At the ND, there are several containment checks. First is a cut on the vertex position, requiring it to be within

---

<sup>1</sup>This sample was enabled by the doubling of our FHC dataset. The corresponding RHC sample was deemed too small for inclusion. Details in Chapter 6.



**Figure 5.1:** Cut flow diagram for the  $\nu_e$  selection at the FD. The rectangles indicate a cut or group of cuts, and the arrows indicate the paths an event can take to potentially end up in one of 3 different samples. The Core, Peripheral, and Low-energy samples each have their own set of event quality and PID cuts. The flow is the same for FHC and RHC, but some cuts are optimized differently for each horn current, and only the FHC Low-energy sample is included in the analysis. The selection process for the ND events follows a similar flow, but with its own cut tunings, and without the nearest slice cut or the Peripheral sample.

a smaller internal volume of the detector known as the fiducial volume, with coordinates

- $-120 \text{ cm} < x < 160 \text{ cm}$
- $-160 \text{ cm} < y < 110 \text{ cm}$
- $150 \text{ cm} < z < 950 \text{ cm}$

The second is a cut requiring the shower be contained in a larger volume, defined by

- $-150 \text{ cm} < x < 170 \text{ cm}$
- $-170 \text{ cm} < y < 150 \text{ cm}$

- $100 \text{ cm} < z < 1240 \text{ cm}$  .

Lastly, the event cannot have any hits in the muon catcher. The tuning metrics used to find the optimal cut values are the signal/background ratio, the Data/MC ratio, the energy resolution, and the fraction of energy escaping the detector [112]. After retuning, the FD cuts from the previous analysis were found to still be optimal, while the ND containment cuts were loosened for this analysis. There were slight increases to the X,Y,Z ranges for the fiducial cut, and a slight increase in the allowed Z range for the shower.

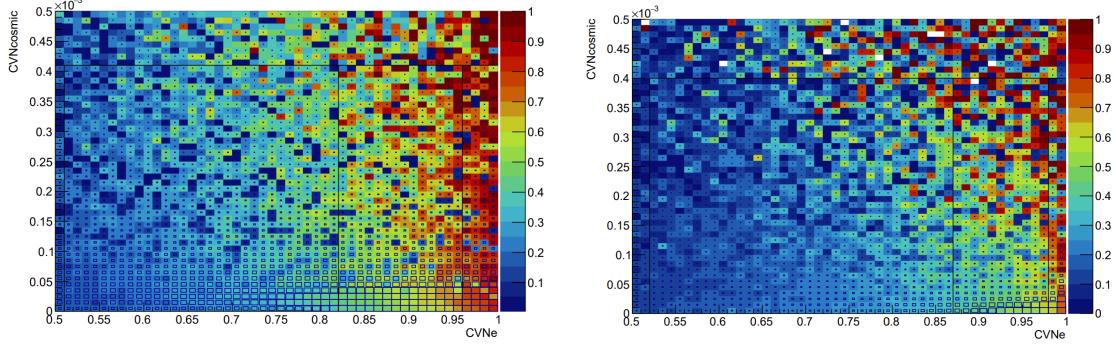
Events passing containment must then pass basic event quality cuts for the core sample. These filter out regions obviously dominated by backgrounds, and select the desired energy ranges for our analysis. In the FD, we require a reconstructed  $\nu_e$  energy between 1 and 4 GeV, between 30 and 150 hits in the event, and a longest prong length between 100 and 500 cm. In the ND we require an energy between 0 and 4.5 GeV, between 20 and 200 hits, and a longest prong between 100 and 500 cm.

Finally, inclusion in the core sample requires passing the particle identification (PID) cuts. These use the output scores from CVN, specifically the electron score, CVNe. We require a CVNe score of at least 0.82 in FHC, and 0.87 in RHC <sup>2</sup>. At the FD, we additionally cut on the cosmic score CVNcos, and the score from the  $\nu_e$  cosmic rejection BDT discussed in Sec. 4.2.4 to mitigate the larger cosmic backgrounds there. We require a CVNcos score less than  $4.1 \times 10^{-4}$  in FHC, and less than  $3.95 \times 10^{-4}$  in RHC. The cosmic rejection BDT score must be at least 0.45 in both FHC and RHC. The ND only requires a cut on CVNe, and uses the same values as the FD. This selects similar samples in each detector, allowing the ND to be used to correct FD predictions (discussed in Sec. 5.4). The optimal cut values are found by maximizing the figure of merit

$$FOM^2 = \frac{S^2}{S + B} \tag{5.1}$$

---

<sup>2</sup>As a reminder, each event's CVN scores add up to 1, with the individual scores of each category indicating the probability that the event is of that type.



**Figure 5.2:** The  $\nu_e$  core sample cut optimization in 3D CoreBDT:CVNcosmic:CVNe space shown as CVNcosmic:CVNe projection for FHC (left) and RHC (right) with the optimal CoreBDT cut of 0.45 applied to both. The color (z-axis) represents signal purity and the box size is proportional to the amount of signal content in each bin. Black lines show optimal cut values, which further divide the sample into the low and high PID bins discussed in Sec. 5.3.1. From [112].

where  $S$  is the number of signal events remaining after the cuts, and  $B$  is the number of background events [112]. All 3 cuts are tuned simultaneously, with separate tunings for FHC and RHC events.

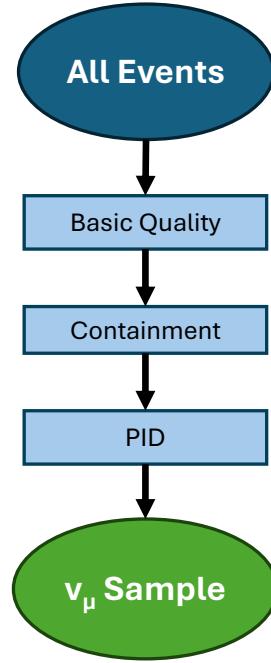
If an FD event failed the containment cut it can still be included in the analysis, through the Peripheral sample. The event quality cuts for the Peripheral sample require a  $\nu_e$  energy between 0 and 4.5 GeV. This wider allowed range reflects the additional uncertainty from energy lost outside the detector. There is an additional requirement that no more than 2.5 GeV can come from the electromagnetic shower. This new requirement was added to remove high-energy beam  $\nu_e$  background events [105]. Equation 5.1 was used to optimize the energy cuts in increments of 0.5 GeV. One of the functions of the containment cut is to remove cosmic backgrounds. The peripheral sample will have higher backgrounds as a result, so we impose stricter PID cuts to account for this. We require a CVNe score of at least 0.82 in FHC, and 0.90 in RHC. The CVNcos score must be less than  $3.5 \times 10^{-5}$  in FHC, and less than  $3.0 \times 10^{-5}$  in RHC. Finally, the peripheral  $\nu_e$  cosmic rejection BDT score must be at least 0.48 in FHC, and 0.52 in RHC. The cuts were tuned using the same method as the core sample described in Sec. 5.2.

If an event passed containment but failed the subsequent core cuts it has one more chance to be included in the analysis, through the low-energy sample. At both the ND and FD, the low-energy sample uses the same event quality cuts as the core sample, but with a different energy cut, requiring a  $\nu_e$  energy between 0.5 and 1.5 GeV. The PID cuts are the same in both detectors. At these lower energies, there are higher backgrounds from cosmic and NC events. Stricter cuts on CVN scores are applied alongside the specialized low-energy BDT described in 4.2.5. For cosmic rejection we require a CVNcos score less than  $3.0 \times 10^{-4}$ , and a  $\nu_e$  cosmic rejection BDT score above 0.47 (this uses the same cosmic BDT as the core sample). To reduce NC backgrounds, we require a CVNnc score less than 0.6. For signal selection, we rely mostly on the low-energy BDT, so the CVNe cut is looser, requiring a score  $> 0.5$ . Much of the low-energy  $\nu_e$  signal occurs at lower CVNe scores as well, so this allows the potential for more  $\nu_e$  events to be reclaimed [103]. Finally, a low-energy BDT score greater than -0.05 is required.

### 5.2.3 $\nu_\mu$ selection

The  $\nu_\mu$  event selection process is more straightforward, with only a single sequence of cuts, outlined in Fig. 5.3. The broad categories are similar to  $\nu_e$ , but optimized for  $\nu_\mu$  events, which amounts to a focus on the reconstructed muon track instead of prongs. We begin with basic reconstruction and event quality cuts. The reconstruction requirements are that the event has at least one 3D Kalman track with a valid ReMID score, a nonzero track energy, and more than 20 hits spanning more than 4 contiguous planes in the detector. The only event quality cut is on the energy, requiring a reconstructed  $\nu_\mu$  energy between 0 and 5 GeV.

The containment cuts at the FD require the event to have no hits within 60 cm of the top, 12 cm of the bottom, 16 cm of the east, 12 cm of the west, and 18 cm from the front or back of the detectors. Additionally, there cannot be any hits in the first 2, or last 3 planes of the detector. Finally, the end point of the track is propagated forward, and the start point of the track is propagated backwards, to determine the number of planes crossed



**Figure 5.3:** Cut flow diagram for the  $\nu_\mu$  selection. Rectangles indicate categories of cuts. There is an identical flow for both ND and FD, as well as FHC and RHC, but with different cut tunings for each. The result is a single sample of  $\nu_\mu$  candidate events for each detector, that can then be binned in different ways described in section 5.3.

before reaching one of the edges of the detector. Both the start and back-propagated tracks must cross at least 6 planes before reaching an edge.

At the ND, we require the hits in the event be contained in a volume defined by

- $-180 \text{ cm} < x < 180 \text{ cm}$
- $-180 \text{ cm} < y < 180 \text{ cm}$
- $-180 \text{ cm} < z < 180 \text{ cm}$

Additionally, the z-coordinate of the start point of the track cannot be greater than 1100 cm. The start and end points of the track are propagated forward/backwards as is done in

the FD, and the forward-propagated track must cross at least 5 planes before encountering an edge, while the back-propagated track must cross at least 10 planes. Both ND and FD cuts were optimized using an identical procedure to the  $\nu_e$  containment cuts [112].

The  $\nu_\mu$  PID cuts use CVNm (the muon score from CVN), the ReMID score from the highest scoring track in the event, and the  $\nu_\mu$  cosmic rejection BDT score. We require a CVNm score above 0.76, a ReMID score above 0.30, and a BDT score above 0.48. All three cuts were simultaneously optimized using the figure of merit Eq. 5.1, with the one difference that for the  $\nu_\mu$  analysis, both right and wrong-sign events count as signal. Several cases were considered: optimizing the FOM across the whole energy range, or optimizing the FOM in the oscillation dip region, for either FHC or RHC modes. The cuts used are those that optimize the FOM in the RHC oscillation dip region. These were found to lower background while keeping a high number of signal events for both maximal and non-maximal mixing values of  $\theta_{23}$  [112].

### 5.3 Analysis Binning

Once the  $\nu_e$  and  $\nu_\mu$  samples have been selected, we bin them in different ways to increase our sensitivity to the oscillation parameters while fitting. This includes splitting the selected samples into various sub-samples, as well as adjusting the number and width of energy bins in the histograms themselves. The binning choices reflect the nature of the sample’s power in determining the oscillation parameters in our analysis. The  $\nu_e$  appearance analysis is more akin to a counting experiment, where the number of  $\nu_e$  vs.  $\bar{\nu}_e$  events is what gives us sensitivity to the mass ordering and value of  $\delta_{\text{CP}}$ . Therefore, the selection process included the additional peripheral sample to reclaim events and the binning choice will focus on enhancing the purity to ensure accurate counts. The  $\nu_\mu$  events give us our sensitivity to  $\Delta m_{32}^2$  and  $\sin^2 \theta_{23}$  through measurement of the location and depth of the oscillation dip between 1-2 GeV. The binning will then focus on the energy resolution, and making a precise measurement of the dip. Binnings are the same between ND and FD selected

samples, to allow for the corrections discussed in subsequent sections.

### 5.3.1 $\nu_e$ Binning

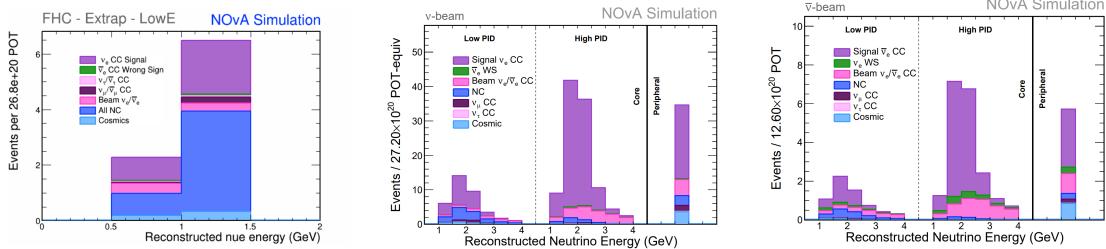
Each of the three  $\nu_e$  samples (core, peripheral, and low-energy) have a unique binning scheme, shown in Fig. 5.4 for the FD. The core sample is split into two sub-samples based on CVNe score. A low-PID, and high-PID sample are defined, with the high-PID requiring a CVNe score of at least 0.96 (0.98) in FHC (RHC), and the low-PID requiring a score between 0.82 (0.87) and 0.96 (0.98) in FHC (RHC). This separation ensures most of the non- $\nu_e$  background (primarily NC, and  $\nu_\mu$  events) is in the low-PID sample. The high-PID sample is then highly pure in  $\nu_e$  events, with the remaining backgrounds coming from the irreducible wrong sign and beam  $\nu_e$  components. The histogram for each subsample uses six 0.5-GeV energy bins across its 1-4 GeV energy range.

The peripheral sample events have missing energy from particles that exited the detector. Therefore, we elect to not use the energy information in our fits, instead using only a single bin to count the number of events. It can also be noted in Fig. 5.4 that the peripheral sample has the highest amount of cosmic background, due to its proximity to the detector edges.

The low-energy sample is already a small sample, so further binning optimization was not pursued, as it is unlikely to lead to gains in sensitivity. We match the histogram binning used in the core sample, with two 0.5-GeV energy bins covering the samples 0.5 - 1.5 GeV energy range.

### 5.3.2 $\nu_\mu$ Binning

The 1 - 2 GeV energy range, where the first oscillation maximum occurs, is the focus of the  $\nu_\mu$  disappearance analysis. We employ a variable binning scheme to improve the precision for measuring the location and depth of the dip in reconstructed energy in that range. Within the dip region we use fine bins of 0.1 GeV, the minimum size allowed by our  $\nu_\mu$  energy resolution. At higher and lower energies we have fewer events and lower

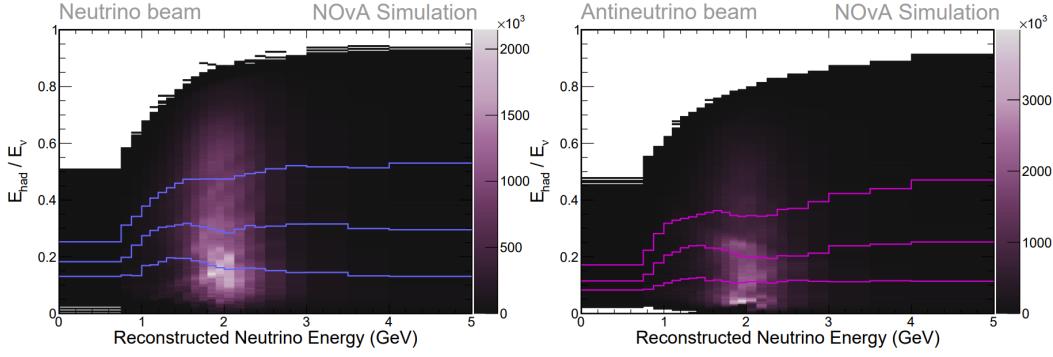


**Figure 5.4:** Predicted FD  $\nu_e$  candidate events for the Low-energy (left), FHC (middle) and RHC (right) samples. The signal events are shown in purple, with the backgrounds stacked below. Predictions were generated with the corrections described in Sec. 5.4, using NOvA’s 2020 analysis best-fit oscillation parameters [113]. From [105].

oscillation probabilities, so we gradually increase the size of the bins there, which reduces the computation time when fitting without sacrificing sensitivity [114]. The number of bins was re-optimized for this analysis, resulting in a modest 1% gain in  $\Delta m_{32}^2$  sensitivity over the previous analysis binning by splitting several bins just outside of the dip region [115].

In Sec. 4.3.2 we described the energy reconstruction for  $\nu_\mu$  events, and saw that the energy resolution for hadronic energy is worse than muon energy, setting the scale for the average FD  $\nu_\mu$  energy resolution at 9.3% (8.0%) in FHC (RHC). To enhance sensitivity to the disappearance measurement, the  $\nu_\mu$  sample is split into sub-samples based on the fraction of hadronic energy present in the event:  $E_{Had}/E_\nu$ . In each bin of reconstructed energy, the FD  $\nu_\mu$  selected sample is split into 4 parts ordered by hadronic energy fraction, each containing 25% of the events in that energy bin. The boundaries defining these quartiles are shown in Fig. 5.5 and separate the  $\nu_\mu$  sample into 4 sub-samples with increasing hadronic energy fraction. This acts as a proxy for binning by the energy resolution, with the lowest hadronic energy quartile having the best energy resolution ( 6.5% in FHC, 5.4% in RHC ), and the highest quartile having the worst resolution ( 12.6% in FHC, 11.2% in RHC ) [106].

In addition to the increased energy resolution in the lower hadronic energy quartiles, we also see most of the backgrounds confined to the highest quartile. This occurs because

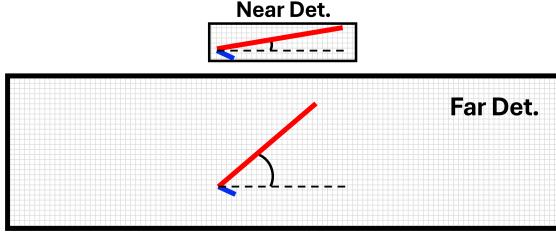


**Figure 5.5:** Hadronic energy fraction vs. reconstructed  $\nu_\mu$  energy for simulated events passing the full  $\nu_\mu$  selection in the FD, for FHC (left) and RHC (right). The z-axis indicates the event count in each bin. The blue and purple lines indicate the four hadronic energy quartile boundaries for each sample. From [112].

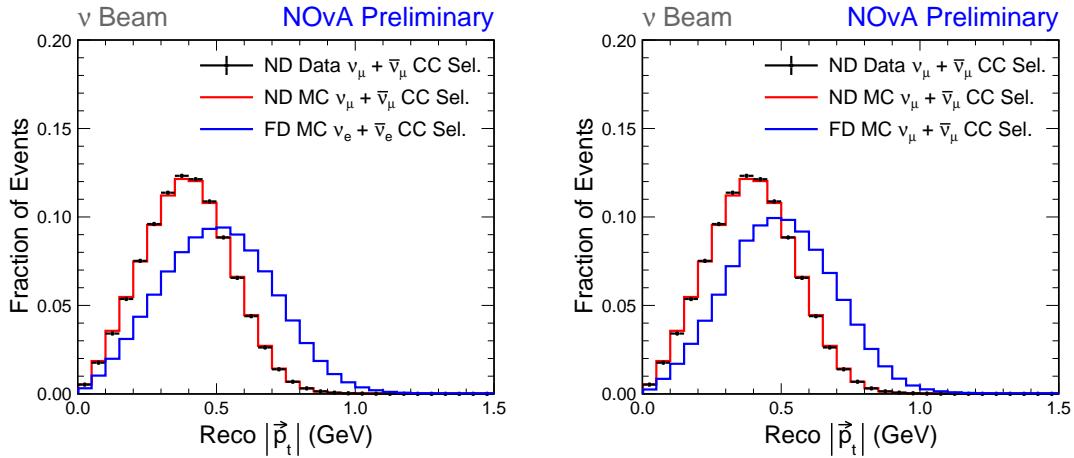
the muon is the primary identifier of a true  $\nu_\mu$  CC event, and less hadronic energy leads to easier muon identification.

### 5.3.3 Transverse Momentum - $p_T$

While the two NOvA detectors are functionally identical, their different sizes lead to different acceptances and selection efficiencies. This is primarily caused by side containment of the muon in  $\nu_\mu$  CC interactions at the ND. Figure 5.6 shows a cartoon of this effect. This leads to a difference in the kinematics of the selected samples, and a sensitivity to systematic uncertainties, particularly cross-section uncertainties [116]. Luckily, there are reconstructed variables that are correlated with this effect, namely the transverse momentum of the interaction,  $p_T$ . Figure 5.7 shows a comparison of  $p_T$  distributions for  $\nu_\mu$  and  $\nu_e$  selected events in the ND and FD. The next section will describe corrections we apply to our FD predictions using ND information. By binning our selected samples in  $p_T$  while applying the corrections, the ND data can constrain any  $p_T$  dependence in our neutrino interaction models, reducing the overall systematic uncertainty in our measurements [116]. The ND  $\nu_\mu$  selected samples are used to define three equally populated bins of  $p_T$  in both FHC and RHC.



**Figure 5.6:** A cartoon showing the different acceptances of the two detectors. A neutrino event with high transverse momentum is more likely to be contained in the far detector.



**Figure 5.7:** Reconstructed transverse momentum distributions. In both plots red is the ND  $\nu_\mu$  selected events. Blue is FD  $\nu_e$  (left), and FD  $\nu_\mu$  (right) selected events. For the FHC beam. From [116].

## 5.4 Constructing Far Detector Predictions

While it is possible to simply use the MC simulation to predict the FD event rates, we can take advantage of the identical designs of the two detectors to apply corrections to the FD predictions using data/MC information from the ND. The result is a reduction in systematic uncertainties that are correlated between detectors, primarily those related to neutrino flux and cross sections. There are two main steps in this process. Each type of neutrino event present in the beam at the ND will propagate differently to the FD. To apply a correction using ND data we first need an estimate of the relative amounts of each beam component in the data. We refer to this estimation process as beam decomposition.

The second step, extrapolation, combines the ND data/MC differences with FD simulation and oscillation parameters to derive a corrected FD prediction. The process is carried out independently for the  $\nu_e$  /  $\nu_\mu$  signal and background channels, which use differing methods for decomposition and extrapolation. Once each is complete, the extrapolated FD event rates for signal and background are recombined, and additional FD-only backgrounds such as cosmics are added in, resulting in the final FD prediction. We will describe the various decomposition and extrapolation methods, and then summarize the techniques used for each analysis sample at the end.

#### 5.4.1 Beam Decomposition

Beam decomposition is used to estimate the relative fractions of each event type in our ND data. In each of the ND selected samples, MC is scaled to exactly agree with the data in each bin of reconstructed energy, and the proportions of each event type can then be adjusted based on additional ND information. The method used to adjust the event proportions depends on the selected sample being decomposed. The  $\nu_\mu$  ( $\bar{\nu}_\mu$ ) selected samples use the signal decomposition method, while the beam backgrounds from the ND  $\nu_e$  selected samples use a combination of the Beam Electron Neutrino (BEN) and Michel decomposition methods (together referred to as “combo decomposition”) or proportional decomposition.

#### Signal Decomposition

The FD signal events for both the  $\nu_e$  appearance and  $\nu_\mu$  disappearance analyses originate from the  $\nu_\mu$  ( $\bar{\nu}_\mu$ ) sample at the ND. The ND  $\nu_\mu$  selected sample is highly pure, and any data/MC discrepancy is assumed to come from the  $\nu_\mu$  ( $\bar{\nu}_\mu$ ) events [105]. Accordingly, the signal decomposition procedure scales only the  $\nu_\mu$  ( $\bar{\nu}_\mu$ ) components to get data/MC agreement. In each bin of reconstructed energy, the simulation and data are compared, and any data/MC disagreement is corrected by adjusting the amount of  $\nu_\mu + \bar{\nu}_\mu$  events so

that the total simulation agrees with data [18]. Any small backgrounds from other event types are not scaled.

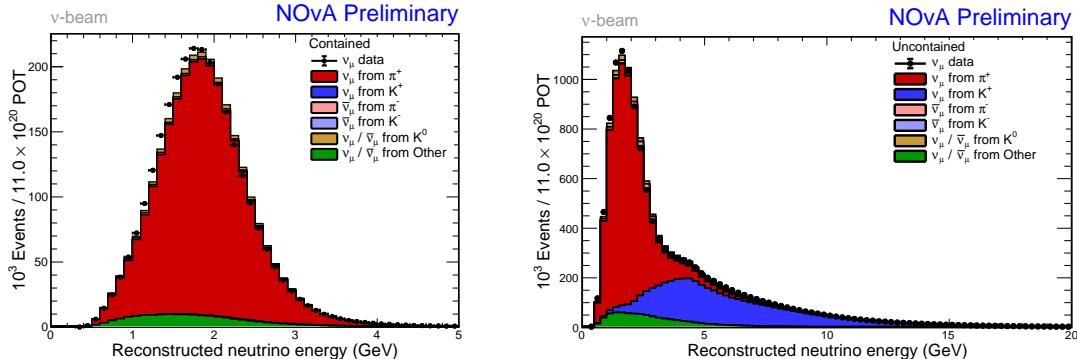
## BEN Decomposition

In the  $\nu_e$  appearance analysis, there are larger FD backgrounds coming from the beam. This includes intrinsic beam  $\nu_e$  CC events and NC events and  $\nu_\mu$  CC events that appear  $\nu_e$ -like. To constrain them, we make use of their corresponding events in the ND  $\nu_e$  selected sample. In FHC, the decomposition process is done in two steps. First, the BEN decomposition method is used to set the proportion of beam  $\nu_e$  CC events in the ND  $\nu_e$  selected sample, using information from the parent mesons of the neutrinos. The Michel decomposition is then used to adjust the remaining ratio of  $\nu_\mu$  CC and NC events, using differences in Michel electron counts in data and MC. In RHC, proportional decomposition is used instead.

As discussed in Sec. 2.2.6, the intrinsic  $\nu_e$  events in the NuMI beam come from kaon decay, as well as the decays of muons coming from pion decay. By constraining the flux of these mesons in the NuMI beam, we can then estimate the amount of  $\nu_e$  events. BEN decomposition does this using samples of contained and uncontained  $\nu_\mu$  events in the ND. The reconstructed  $\nu_\mu$  energy distributions of these samples are shown in Fig. 5.8, with their parent meson ancestry highlighted. The selection for the contained sample is the ND  $\nu_\mu$  selection from 5.2.3, and the selection for the uncontained sample is the same but excluding the containment cut [117]. The contained  $\nu_\mu$  sample is originating almost entirely from pion decays, while the uncontained sample has a large population of events from kaon decay. First, the pion rate is identified using the contained  $\nu_\mu$  sample. A set of weights  $\omega_{\nu_\mu}$  are calculated by taking the ratio of the estimated number of  $\nu_\mu$  CC events from pions and kaons in data vs. MC. These ratios are propagated back to the parent pions as a function of pion momentum, to get a weight for the pions as a function of momentum ( $p_T, p_z$ ) [16]. The kinematics of the pions that produced the  $\nu_\mu$  events are very similar to those of the pions that produce the  $\nu_e$  events, as shown in Fig. 5.9. Therefore, we can apply the same

pion event weights calculated from the  $\nu_\mu$  sample to the events with pion ancestry in the selected  $\nu_e$  sample. The corrected rate of  $\nu_e$  events from pions is then found by summing over pion ( $p_T, p_z$ ) in each reconstructed energy bin [118].

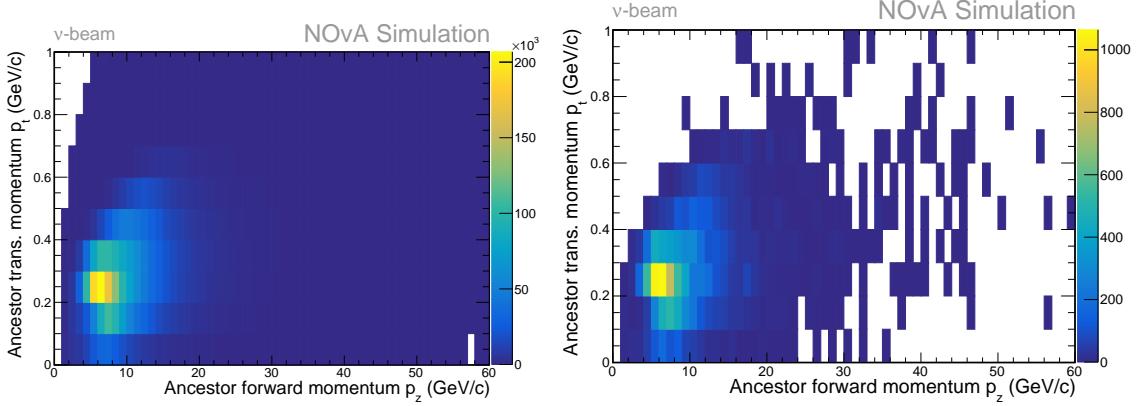
The remaining beam  $\nu_e$  events from kaon decays are corrected using the uncontained  $\nu_\mu$  sample. A kaon scale,  $S_K$  is calculated by computing a ratio of the number of selected  $\nu_\mu$  CC events from kaons in data vs. MC. First, the pion component of the uncontained  $\nu_\mu$  sample is corrected using the pion weights from the contained sample. Then, the bins in the 4.5-10 GeV range are combined to account for energy loss in the uncontained sample, and the kaon component is adjusted to achieve data/MC agreement [18]. This scaling is then applied to the ND  $\nu_e$  selected events originating from kaons.



**Figure 5.8:** Reconstructed  $\nu_\mu$  energy distributions for the FHC contained (left) and uncontained (right)  $\nu_\mu$  selected samples in the ND. The events are split by their corresponding parent meson, and used to constrain the pion and kaon flux in the BEN decomposition. From [41].

### Michel Decomposition

Once the proportion of  $\nu_e$  events in the ND  $\nu_e$  selected sample has been set by the BEN decomposition, the remaining ratio of  $\nu_\mu$  CC and NC events is determined by the Michel decomposition. This is done using differences in the rates of Michel electrons between the two types of event. While both types of events can produce Michel electrons through the hadronic system (primarily via muons from pion decays), the presence of the primary muon in the  $\nu_\mu$  CC events means that on average the  $\nu_\mu$  CC events will contain one more



**Figure 5.9:** Simulated transverse vs. forward momentum for parent pions of  $\nu_\mu$  (left) and  $\nu_e$  (right) CC events in the ND, for FHC mode. From [41].

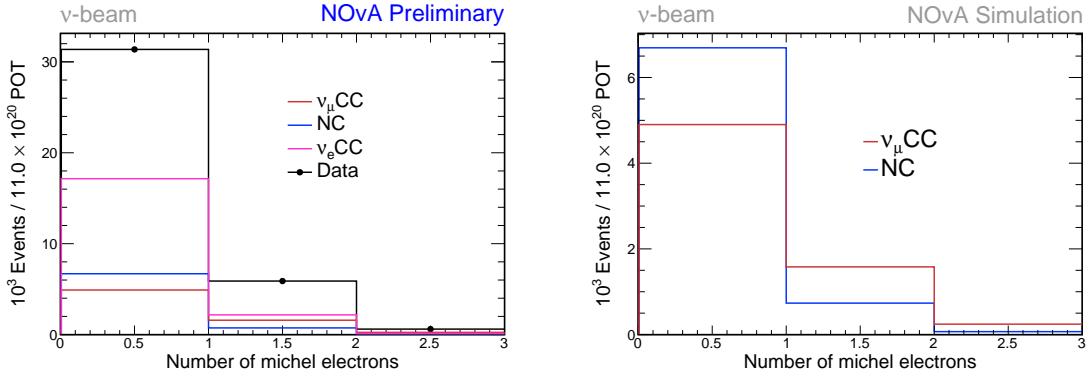
reconstructed Michel electron than the NC events. Figure 5.10 shows the distribution of the number of reconstructed Michels per event, for the ND  $\nu_e$  selected sample.

The decomposition procedure is carried out separately in each bin of reconstructed energy for the ND  $\nu_e$  selected sample. The number of  $\nu_e$  events in each bin is set by the BEN decomposition, and kept fixed. The number of  $\nu_\mu$  CC and NC events is then scaled up so that data/MC agree in each bin. We then construct separate NMichel distributions analogous to Figure 5.10 for each bin of reconstructed energy. Finally, the ratio of  $\nu_\mu$  CC to NC events are adjusted in each energy bin by performing a log-likelihood fit to data using the NMichel distribution for that bin [119].

The Michel decomposition is only performed for energy bins in which at least 20% of events in that bin have a reconstructed Michel electron. This is to ensure an accurate fit can be performed to the NMichel distribution [119]. For bins without enough Michels, a fallback to the proportional decomposition method is used (described below).

### Proportional Decomposition

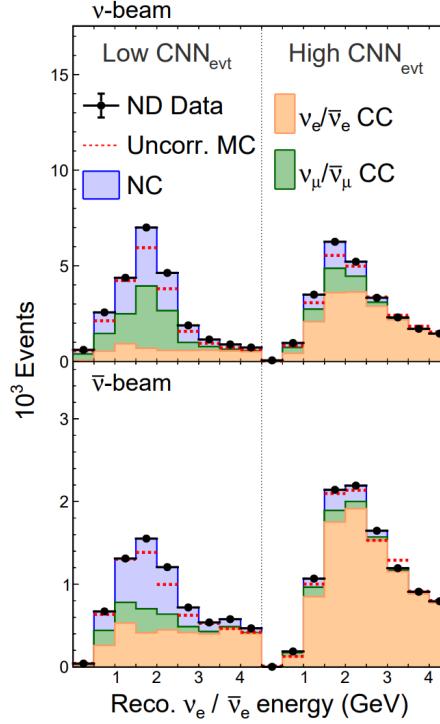
Proportional decomposition is a simpler method which only scales the MC to agree with data in each bin of reconstructed energy, without further adjustments to the proportions of each event type. It assumes the simulation has correctly modeled the relative ratios of



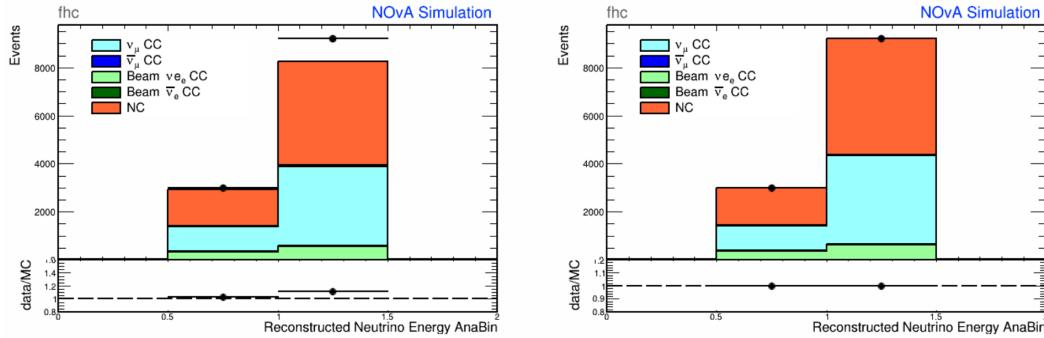
**Figure 5.10:** The number of reconstructed Michels per event in the ND  $\nu_e$  selected sample, before corrections are applied, for all  $\nu_e$  energies. Left shows a data/MC comparison, and right spotlights the simulated  $\nu_\mu$  CC and NC events from the left plot, which are used in the Michel decomposition. These Michels come from MEFinder (section 4.1.6) and must pass an MID and  $\Delta T$  cut. From [41].

the event types, but not the overall normalization in each reconstructed energy bin.

Currently the BEN+Michel decomposition method is used for the FHC Core  $\nu_e$  backgrounds, and proportional decomposition is used for the RHC  $\nu_e$  backgrounds, the FHC low-energy  $\nu_e$  backgrounds, and as a fallback for low-statistics bins in the FHC Michel decomposition. A combo decomposition method was developed for RHC, but requires storing additional ND RHC data/MC information in the files, greatly adding to the computation cost [105]. For this reason the proportional decomposition method was chosen instead. Figure 5.11 shows the ND Core  $\nu_e$  selected sample before/after decomposition has been applied, along with the relative scalings applied by the FHC combo decomposition. Figure 5.12 shows the ND low-energy  $\nu_e$  selected sample before and after the proportional decomposition is applied. Note that there is no ND sample corresponding to the FD Peripheral sample. For the FD peripheral sample backgrounds we will use the decomposition from the high-PID Core  $\nu_e$  bins, since they share similarly strict PID requirements.



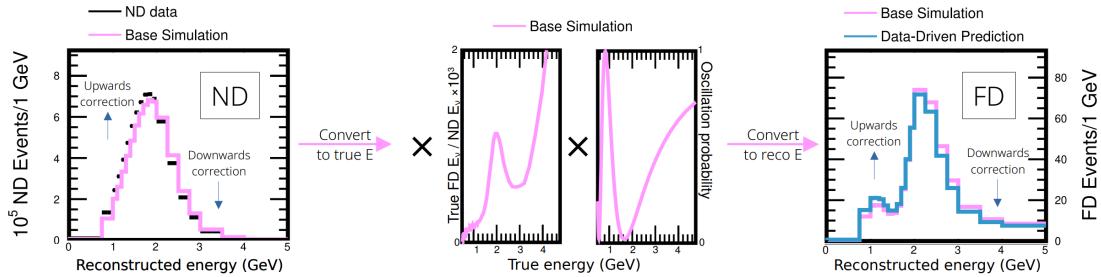
**Figure 5.11:** The ND  $\nu_e$  selected samples in FHC (top) and RHC (bottom) after their corresponding decomposition methods have been applied. The total MC before the correction is shown as a dotted red line. After correction, the MC agrees exactly with the data, by construction. From [105].



**Figure 5.12:** The ND low-energy  $\nu_e$  selected sample, before and after applying proportional decomposition. These are the events that will oscillate to become the beam backgrounds to the FD low-energy  $\nu_e$  selected sample. The first bin had a scaling of 2.56%, and the second bin had a scaling of 11.7% applied. From [120].

### 5.4.2 Extrapolation

Extrapolation is performed for both signal channels, and for the large  $\nu_e$  backgrounds.  $\nu_e$  backgrounds from other oscillation channels, and all  $\nu_\mu$  backgrounds, are too small for data-driven corrections to be effective, and so are simply taken directly from the FD simulation. A schematic of the extrapolation procedure is shown in Fig. 5.13, using the  $\nu_\mu \rightarrow \nu_\mu$  signal extrapolation as an example.



**Figure 5.13:** Schematic diagram showing the extrapolation procedure for the  $\nu_\mu \rightarrow \nu_\mu$  signal events. From [29].

The first step is the decomposition described above, where the ND data is decomposed using the corrected ND MC. Once we have a corrected ND spectrum, we can convert it to true energy before extrapolating. This is done via a reconstructed-to-true energy smearing matrix derived from ND simulation [121]. This requires an accurate conversion between reconstructed and true energy, and so extrapolation in bins of true energy can only be done for components that have good energy resolution, i.e., the  $\nu_\mu \rightarrow \nu_\mu$  and  $\nu_\mu \rightarrow \nu_e$  signal components (referred to as “truth extrapolation”) [121]. The backgrounds contain mis-identified events, and so we do not expect the energy estimators to perform well. For the background components, the corrected ND spectrum is not converted to true energy, and the extrapolation is done in bins of reconstructed energy (referred to as “reco extrapolation”).

The next step multiplies the decomposed ND data by a ratio of the uncorrected Far/Near simulation, to get a FD prediction in true energy. For truth extrapolation this

takes the form

$$F_{\alpha \rightarrow \beta}^{Pred}(E_i^T, B_j) = \frac{N_{\alpha}^{Data}(E_i^T) \times F_{\alpha \rightarrow \beta}^{MC}(E_i^T, B_j)}{N_{\alpha}^{MC}(E_i^T)} \quad (5.2)$$

where  $N$  and  $F$  represent event rates in the near and far detectors, normalized to data POT in the appropriate detector.  $MC$  indicates a rate derived directly from simulation,  $Data$  is a rate from the decomposition, and  $Pred$  is a predicted rate calculated by the extrapolation. These rates are calculated for neutrino flavors  $\alpha, \beta$  in the ND, FD respectively <sup>3</sup>. Terms in parenthesis indicate a binning, in either true energy  $E^T$  indexed by  $i$ , or analysis binning  $B$  indexed by  $j$  (nominally the reconstructed energy, but which can take on any reconstructed variable) [121].

For reco extrapolation the equation takes the form

$$F_{\alpha \rightarrow \beta}^{Pred}(E_i^T, B_j) = \frac{N_{\alpha}^{Data}(B_j) \times F_{\alpha \rightarrow \beta}^{MC}(E_i^T, B_j)}{N_{\alpha}^{MC}(B_j)} \quad (5.3)$$

Where the only difference is the binning of ND samples in reconstructed energy, since the reco-to-true matrix was not applied.

This can then be multiplied by the oscillation probability, and converted back from true energy to the analysis bins using the corresponding FD true-to-reco energy smearing matrix.

$$F_{\alpha \rightarrow \beta}^{Pred}(B_j, \vec{\theta}) = \sum_i F_{\alpha \rightarrow \beta}^{Pred}(E_i^T, B_j) \times P(E_i^T, \vec{\theta}) \quad (5.4)$$

Here, the oscillation probability  $P$  is for a particular set of oscillation parameters  $\vec{\theta}$ . Factoring the prediction in this way allows us to easily recompute the rates at many different values of the oscillation parameters ( for example, during fitting), without recomputing all of the  $F_{\alpha \rightarrow \beta}^{Pred}(E_i^T, B_j)$  terms [121].

There are two equivalent ways to look at this application of the extrapolation correction,

---

<sup>3</sup>For the purposes of extrapolation, neutral current is considered as a separate, non-oscillating flavor. Additionally, there are vanishingly few  $\nu_{\tau}$  in the near detector (ND MC does not include them at all), so the allowed flavors are  $\alpha \in \{\nu_e, \bar{\nu}_e, \nu_{\mu}, \bar{\nu}_{\mu}, NC\}$  and  $\beta \in \{\nu_e, \bar{\nu}_e, \nu_{\mu}, \bar{\nu}_{\mu}, \nu_{\tau}, \bar{\nu}_{\tau}, NC\}$  [121].

which in simpler terms takes the form

$$F_{Pred} = N_{Data} \times \frac{F_{MC}}{N_{MC}} = \frac{N_{Data}}{N_{MC}} \times F_{MC} . \quad (5.5)$$

The first expression emphasizes a reweighting of ND data by the simulated far/near ratio, which leads to the reduction in systematic uncertainties correlated between detectors. The second expression emphasizes a reweighting of FD simulation by a data/MC correction. The extrapolation captures both effects, so it is helpful to think of both reweightings simultaneously.

Below is a summary of the different decomposition and extrapolation techniques used for each of the primary samples in the analysis.

### FD $\nu_\mu \rightarrow \nu_\mu$ Signal

Extrapolates the signal events from the ND  $\nu_\mu$  selected sample to the FD  $\nu_\mu$  selected sample, using signal decomposition and truth extrapolation. During extrapolation, the samples are split into four bins of hadronic energy, and three bins of  $p_T$ , as described in section 5.3, for a total of 12 extrapolated samples each in FHC and RHC. The  $p_T$  bins are summed together at the end, and the 4 FD hadronic energy samples are carried forward in the analysis.

### FD $\nu_\mu \rightarrow \nu_\mu$ Background

All beam backgrounds in the  $\nu_\mu$  analysis are small, and taken directly from the FD simulation without corrections. Events are split into their corresponding hadronic energy bins.

### FD $\nu_\mu \rightarrow \nu_e$ Signal

Extrapolates the ND  $\nu_\mu$  selected sample to the FD  $\nu_e$  selected sample, using signal decomposition and truth extrapolation. The samples are split into three bins of  $p_T$  for the

extrapolation, and recombined into their FD analysis bins at the end.

### FD $\nu_\mu \rightarrow \nu_e$ Background

For the large  $\nu_e$  backgrounds (beam  $\nu_e$ , NC,  $\nu_e$ -like  $\nu_\mu$  CC) we extrapolate the ND  $\nu_e$  selected samples to the FD  $\nu_e$  selected samples, using combo decomposition for the FHC core and peripheral samples, and proportional decomposition for the low-energy sample and RHC core and peripheral samples. Note that the peripheral sample uses the ND high-PID bins from the core sample for its decomposition, as there is no corresponding ND peripheral sample. Reco extrapolation is used in all cases, without  $p_T$  binning. All other beam backgrounds in the  $\nu_e$  analysis are small, and taken directly from the FD simulation without corrections.

### 5.4.3 Cosmic and Rock Backgrounds

As mentioned in the introduction, the final step in constructing the FD predictions is to incorporate additional backgrounds coming from cosmic events, and beam neutrino interactions in the rock surrounding the detector. As discussed in Secs. 2.3 and 3.1.3, the rock sample comes from a special dedicated simulation of beam neutrino events outside the detector, and the cosmic sample uses real data taken from the FD cosmic trigger to avoid a costly simulation. As a reminder, this data is taken outside the NuMI spill window so there is no potential overlap with beam events.

For both backgrounds we apply the full  $\nu_\mu$  and  $\nu_e$  selection cuts and analysis binnings to get the corresponding backgrounds for each selected sample. The rock prediction is taken directly from the simulation, without extrapolation applied, and can be scaled to match the POT of the data used in the analysis. The cosmic sample comes from data outside the beam spill, which has no associated POT. Instead, this sample is scaled based on time, scaling down the cosmic data livetime to match the beam data's livetime. For the  $\nu_\mu$  samples, the cosmic distributions are also smoothed to get a better estimate of the energy dependence (described in Sec. 5.5.4). The backgrounds for each sample are then

summed into the final predictions. The total cosmic background events before and after scaling are summarized in Table 5.1. The number of predicted rock background events are summarized in Table 5.2.

FHC		RHC	
	no scaling	885.5s livetime	no scaling
$\nu_e$ All	2047	5.40	1047
$\nu_\mu$ Q1	223	0.49	81
$\nu_\mu$ Q2	237	0.52	99
$\nu_\mu$ Q3	352	0.77	151
$\nu_\mu$ Q4	1235	2.69	467

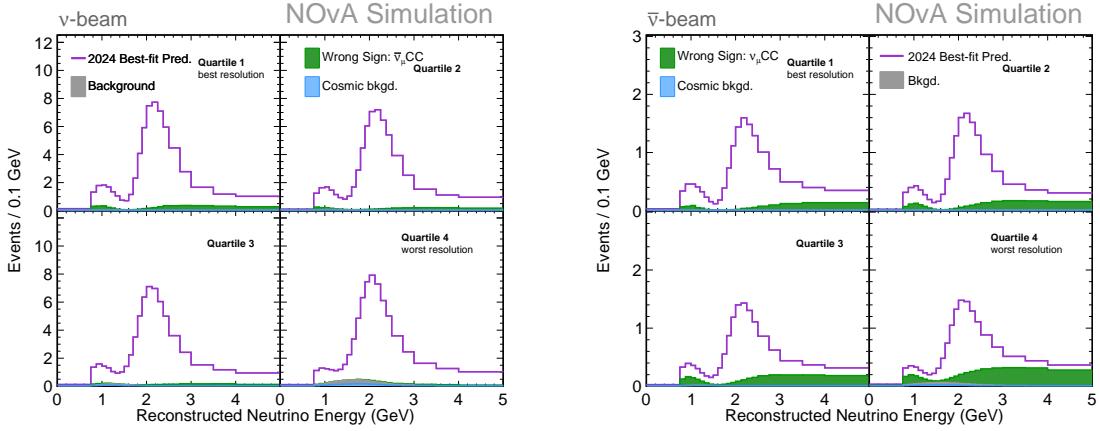
**Table 5.1:** Cosmic background events selected for all  $\nu_e$  samples, and each  $\nu_\mu$  quartile, before and after scaling to match beam livetimes. Before scaling, the FHC cosmic sample had a livetime of 406,368.56 seconds, and the RHC a livetime of 367,915.41 seconds. From [122].

	FHC	RHC
$\nu_e$	2.62	0.38
Low-energy	0.15	-
$\nu_\mu$	0.020	0.005

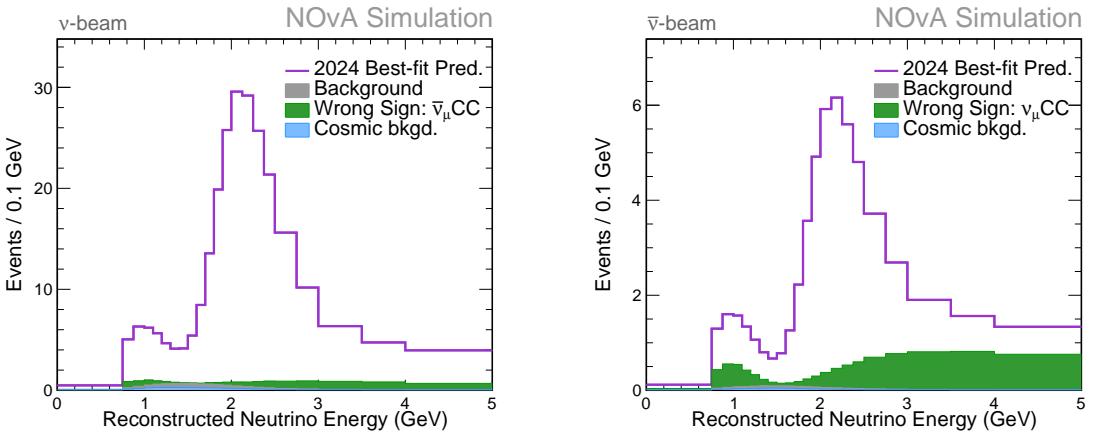
**Table 5.2:** Predicted number of additional events from rock interactions in the FD selected samples. From [105].

#### 5.4.4 Final FD Predictions

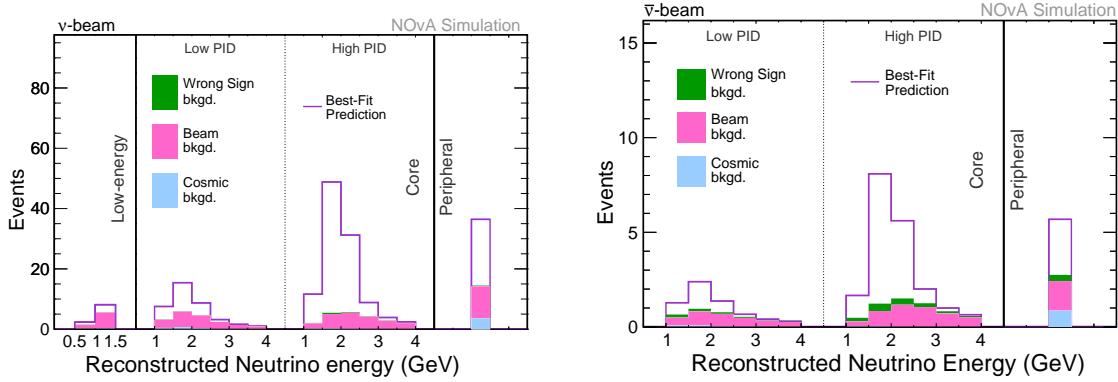
The final FD predictions for the  $\nu_\mu$  FHC and RHC sample in hadronic energy bins are shown in Fig. 5.14. The prediction with all hadronic energy bins summed together is shown in Fig. 5.15. The  $\nu_e$  predicted samples are shown in Fig. 5.16. The total predicted event counts are summarized in Table 5.3.



**Figure 5.14:** Predictions for the FD  $\nu_\mu$  selected events in the FHC (left) and RHC (right) samples, split by hadronic energy fraction, with extrapolation applied. The total prediction is shown in purple. The background components are shown as a stacked histogram below this, with the remaining whitespace indicating the signal component. Predictions were generated using the best-fit oscillation parameters found in Sec. 7.2.



**Figure 5.15:** Predictions for the FD  $\nu_\mu$  selected events in the FHC (left) and RHC (right) samples, with all hadronic energy bins summed, and with extrapolation applied. The total prediction is shown in purple. The background components are shown as a stacked histogram below this, with the remaining whitespace indicating the signal component. Predictions were generated using the best-fit oscillation parameters found in Sec. 7.2.



**Figure 5.16:** Predicted FD  $\nu_e$  candidate events for the FHC (left) and RHC (right) samples, with extrapolation applied. The total prediction is shown in purple. The background components are shown as a stacked histogram below this, with the remaining whitespace indicating the signal component. Predictions were generated using the best-fit oscillation parameters found in Sec. 7.2.

Sample	FHC			RHC	
	$\nu_\mu$	$\nu_e$	Low-energy	$\bar{\nu}_\mu$	$\bar{\nu}_e$
$\nu_\mu \rightarrow \nu_\mu$	372.3	4.3	0.3	24.4	0.2
$\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$	24.5	0.1	0.0	71.5	0.2
$\nu_\mu \rightarrow \nu_e$	0.4	125.3	3.4	0.0	2.1
$\bar{\nu}_\mu \rightarrow \bar{\nu}_e$	0.0	1.8	0.1	0.0	18.9
Beam $\nu_e + \bar{\nu}_e$	0.1	26.1	0.8	0.0	6.5
NC	5.5	16.8	5.3	0.8	2.0
Cosmic	4.4	5.5	0.5	0.7	1.1
Others	1.5	0.8	0.1	0.2	0.1
Signal	397.6	125.3	3.4	96.0	18.9
Background	11.0	55.4	7.1	1.7	12.2
Total	408.6	180.7	10.5	97.7	31.1

**Table 5.3:** Predicted event counts for the  $\nu_\mu$ ,  $\nu_e$  and Low-energy selected samples in the neutrino beam, and  $\bar{\nu}_\mu$  and  $\bar{\nu}_e$  samples in the antineutrino beam. The low/high PID, and Peripheral samples are combined in the  $\nu_e$  ( $\bar{\nu}_e$ ) columns, while the low-energy sample is shown separately due to its novel status. Signal in the  $\nu_\mu$  ( $\bar{\nu}_\mu$ ) columns includes wrong-sign events and some  $\nu_\mu$  from Others. Predictions were generated using the best-fit oscillation parameters found in Sec. 7.2.

## 5.5 Systematic Uncertainties

When extracting oscillation parameters through our fits to data, we account for over 60 systematic uncertainties covering many aspects of our simulated model [105]. Described in Sec. 5.6, likelihoods are minimized with respect to systematic terms as well as oscillation parameters. To determine the likelihood as a function of systematics, a new set of simulated predictions are created for each systematic, where the effect has been applied at the  $\pm 1$  and  $\pm 2\sigma$  levels. This is done either by using functions to reweight the energy distributions of the nominal predictions, or for more complex effects, by making new predictions from scratch, rerunning the full reconstruction chain over dedicated samples with a systematically-shifted simulation. Within the fits, a continuous systematic range is necessary, so the final spectra from the  $1\sigma$ , and  $2\sigma$  predictions are interpolated to get values between the nominal and  $1\sigma$ , and  $2\sigma$  levels [111]. The broad categories of systematic are flux, neutrino cross section, detector response and calibration, and other systematics.

### 5.5.1 Flux Systematics

The flux systematic covers uncertainty in our simulation coming from mismodeling of the production and transport of hadrons in our beamline (described in Sec. 3.1.1), which subsequently decay to create neutrinos. The Package to Predict the Flux (PPFX) [55], which provides data-driven corrections to our proton target cross-section, is also used to construct the uncertainty using a multiverse technique. A statistical ensemble of 100 “universes” are considered where the cross sections in our proton target have been varied within the allowed uncertainties of the PPFX data constraints [123]. For each of these universes, an additional 20 scenarios are considered that vary parameters related to the subsequent transport of hadrons, such as focusing horn position and current, beam position and spot size, and target position [18]. The result is 2000 fluctuated universes, each with their own weights that can be used to correct the default simulation of neutrino interactions to achieve that universe [124]. For each universe, a covariance matrix is constructed in

true energy bins of the resulting neutrino interactions. These matrices are then summed, forming an average covariance matrix. Since the flux uncertainties are known to be highly correlated across true energy bins since we are off-axis, and to reduce the dimensionality of our fits, we can summarize their effects using a Principal Component Analysis (PCA) on the averaged covariance matrix [124]. The matrix is diagonalized, with the eigenvectors taken as the principal components (PCs) and the size of their eigenvalues determining their ranking. The 5 largest PCs are chosen as the final flux systematics, and scaled up by 25% to ensure proper coverage of the full underlying systematic effect [105].

### 5.5.2 Cross-Section Systematics

Cross sections and final state interactions (FSI) represent the largest category of systematics, with 78 systematic “knobs” that modify interaction models in the fit [105]. This is done through GENIE [52], which has weights that can be applied to individual events using their neutrino truth information [40]. There are knobs affecting all of the simulated interaction modes described in Sec. 3.1.2. New for this analysis are additional knobs targeting RES and DIS events, to account for mismodeling of pion production in the transition region between the two types of event [125].

Since the extrapolation process reduces the impact of cross section uncertainties, many of the systematics will have a small effect on the final fit results. To reduce computation time when fitting, we split the list into “large” and “small” groups, with the large systematics treated individually, and the small systematics encompassed in a PCA. The distinction is determined by the  $\Delta\chi^2$  the  $+/-1\sigma$  shifted samples produce with respect to the nominal predictions, with any that have an impact above  $\Delta\chi^2 = 0.005$  considered as large [126]. Additionally, any systematics that are being included for the first time are placed in the “large” group regardless of  $\chi^2$ . Thirty of the systematics fall in the large category, summarized in Table 5.4.

We conduct a PCA on the remaining 48 small systematic knobs to summarize their effect. Similar to the flux systematic, we use a multiverse technique where 1000 universes

are generated with variations in the systematics. The averaged covariance matrix in true neutrino energy bins is constructed, and diagonalized, with the eigenvectors of the matrix taken as the PCs [126]. The eigenvalues are used to rank the PCs, with the 8 largest included as individual systematics in the analysis, with a 35% upscaling to ensure good coverage of all the underlying effects [105].

**Table 5.4:** Name and short description of the effect in simulation for the 30 “large” cross section systematic knobs. Those with a star are new for this analysis. From [105, 111, 125].

Systematic Name	Effect
ZNormCCQE	Normalization parameter in the CCQE z-expansion axial form factor
ZExpAxialFFSyst2020_EV{1,2,3,4}	Four correlated CCQE z-expansion axial vector shape variations
MECEnuShape2020	$E_\nu$ dependence of MEC for neutrinos
MECEnuShape2020AntiNu	$E_{\bar{\nu}}$ dependence of MEC for antineutrinos
MECShape2024Nu	$(q_0,  \vec{q} )$ dependence of MEC for neutrinos
MECShape2024AntiNu	$(q_0,  \vec{q} )$ dependence of MEC for antineutrinos
MECInitStateNPFrac2020Nu	Fraction of MEC interactions on neutron-proton pairs for neutrinos
MECInitStateNPFrac2020AntiNu	Fraction of MEC interactions on neutron-proton pairs for antineutrinos
MaCCRES	Mass parameters in the axial form factors for resonant production in CC events
MaNCRES	Mass parameters in the axial form factors for resonant production in NC events
MvCCRES	Mass parameters in the vector form factors for resonant production in CC events
MvNCRES	Mass parameters in the vector form factors for resonant production in NC events

**Table 5.4:** Name and short description of the effect in simulation for the 30 “large” cross section systematic knobs. Those with a star are new for this analysis. From [105, 111, 125].

RPAShapeenh2020	Higher- $Q^2$ Enhancement for Random Phase Approximation in CCQE
RPAShapesupp2020	Low- $Q^2$ Suppression for Random Phase Approximation in CCQE
LowQ2RESsupp2020	Low- $Q^2$ Suppression in RES events
RESvpvnRatioNuXSecSyst *	Relative RES cross-section scaling of $\frac{\sigma(\nu+p)}{\sigma(\nu+n)}$
RESvpvnRatioNubarXSecSyst *	Relative RES cross-section scaling of $\frac{\sigma(\bar{\nu}+p)}{\sigma(\nu+n)}$
RESDeltaScaleSyst *	Scales the normalization of RES interactions which produce a $\Delta$
RESOtherScaleSyst *	Scales the normalization of RES interactions producing higher order resonances
DISNuHadronQ1Syst *	Relative scaling of $\nu$ DIS interactions with 2 final state hadrons with total $Q=1$
DISNuBarHadronQ0Syst *	Relative scaling of $\bar{\nu}$ DIS interactions with 2 final state hadrons with total $Q=0$
DISvnCC1pi_2020	Normalization factor of $1\pi$ final states in DIS scattering of neutrinos from neutrons
hNFSI_MFP_2024	Neutrino mean free path in FSI
hNFSI_FateFracEV1_2024	Largest of three FSI correlated ‘fate fraction’ shifts
radcorrnu	$\nu_e/\nu_\mu$ cross-section differences from radiative corrections
radcorrnuubar	$\bar{\nu}_e/\bar{\nu}_\mu$ cross-section differences from radiative corrections
2ndclasscurr	Cross-section differences from second-class currents

### 5.5.3 Detector Response and Calibration Systematics

Some of the most significant uncertainties are related to the model of our detector response and the calibration process. Modifying the underlying parameters for these systematics

can have many effects that are hard to quantify with a simple scaling of neutrino energy. For example, changing the amount or energy of hits, which then alters the reconstruction and PID performance. Therefore, to quantify their effects necessitates a partial or full re-simulation of the ND and FD samples. For each of the  $\pm 1$  and/or  $\pm 2\sigma$  shifts, a new sample is made. Some systematics consider all four shifts, while some only consider  $\pm 1\sigma$ , or are one-sided shifts in the + or - directions. These samples are processed through the full reconstruction chain, and separate predictions are generated for each shift of each systematic. These are described below.

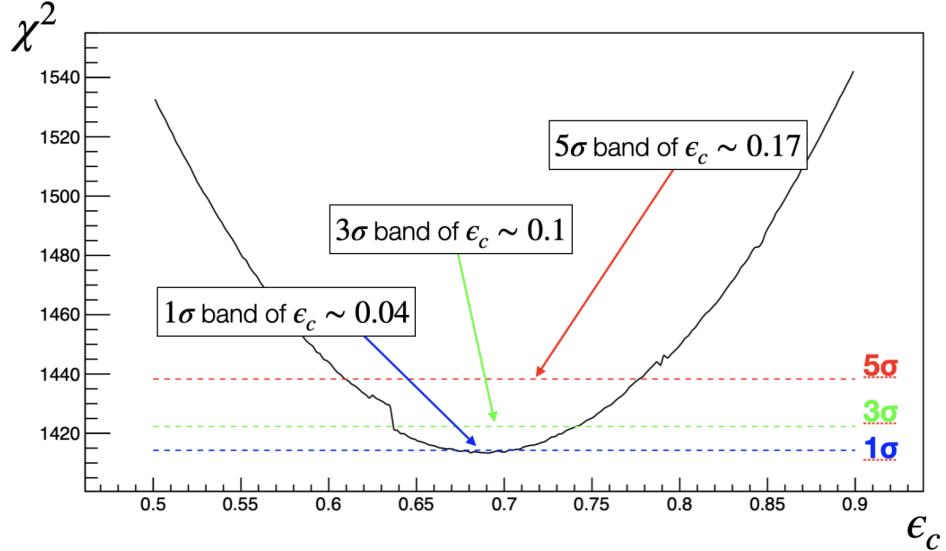
### Light Level and Cherenkov

The light level and Cherenkov systematics modify parameters in the light level tuning procedure described in Sec. 3.1.2 to get  $\pm 1\sigma$  shifted predictions. Specifically, the light level systematic varies the  $F_{\text{view}}$  detector-specific scaling parameters in equation 3.1, while the Cherenkov uncertainty varies the Cherenkov scaling parameter  $\epsilon_C$ . The amount that each parameter needs to be adjusted up/down to correspond to a  $\pm 1\sigma$  shifted prediction is determined during the light-level tuning procedure<sup>4</sup>. For  $\epsilon_C$ , the amount is found by profiling the parameter during the light level tune fit, while allowing the attenuation and FD view scale parameters to float (the ND parameters are kept fixed, and are covered by the light level systematic) [127]. Figure 5.17 shows the resulting profile, with  $1\sigma$ ,  $3\sigma$ , and  $5\sigma$  ranges of the parameter indicated. The  $3\sigma$  range was conservatively chosen for the systematic, which shifts the value of  $\epsilon_C$  by  $\pm 0.05$  ( $\pm 6.2\%$ ) [127].

For the light level systematic, a similar approach was used, where the ND and FD  $F_{\text{view}}$  scaling parameters were profiled and the  $3\sigma$  confidence interval was used to set the scale of the uncertainty. For each detector, a likelihood surface was constructed, with the X and Y-view scaling parameters as the axes, and the other parameters profiled as nuisance parameters in the fit [127]. This means that at each point in the surface the X and Y-view

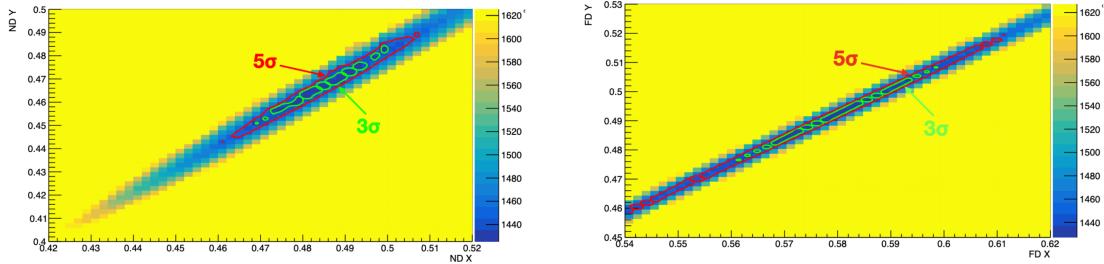
---

<sup>4</sup>As a reminder, the light tune procedure is a joint fit which varies the light level parameters to achieve data/MC agreement in several selected samples simultaneously (ND muons, protons, and cosmics, and FD cosmics).



**Figure 5.17:** Profile of the Cherenkov scaling parameter in the light level tune fit. ND view scaling parameters were held fixed, while all other parameters were allowed to float. The  $3\sigma$  range here was used to set the scale of the Cherenkov uncertainty. From [127].

parameters are set to their axis values, and the other parameters are fit to the data to find the values that minimize the  $\chi^2$  for that point. The resulting surfaces are shown in Figure 5.18. For each surface, the width of the  $3\sigma$  ellipse covers a range of approximately  $\pm 0.05$  on each axis. Therefore, the light-level uncertainty was set as a  $\pm 5\%$  scaling of the  $F_{\text{view}}$  parameters for each detector. Therefore the  $\pm 1\sigma$  shifted systematic predictions will have an increase/decrease in the light level of 5%.



**Figure 5.18:** The ND (left) and FD (right) profiles of the X and Y  $F_{\text{view}}$  scaling parameters in the light level tune fit. 3 and  $5\sigma$  bands are drawn, and the extent of the  $3\sigma$  bands were used to set the light level uncertainty. From [127].

## Calibration Energy Scale

The absolute calibration process uses stopping cosmic ray muons as a standard candle to set the absolute energy scale (Sec. 3.2.2). ND Data/MC ratios of additional standard candles such as beam muon and proton  $\frac{dE}{dx}$ , neutral pion decays, and Michel electrons are used to define the uncertainty in the energy scale. Historically, NOvA used a 5% uncertainty set by the largest discrepancy, the beam muon and proton  $\frac{dE}{dx}$ . For this analysis, the updated light level tune resulted in significant improvement to muon and proton  $\frac{dE}{dx}$ , and all standard candles now fall under 2% uncertainty [127]. However, it was decided to maintain the 5% uncertainty level until more studies could be conducted that consider different contributions from electromagnetic, hadronic, and muonic systems [127].

Since we cannot effectively determine whether the ND data/MC ratios extend to the FD, we consider two separate systematics for the energy scale: a fully correlated and anti-correlated uncertainty, to bracket all possible scenarios. The fully correlated systematic applies the same +/-5% shifts to both ND and FD. The anti-correlated systematic applies opposite +5% and -5% shifts in the ND and FD [127].

## Calibration Shape

The calibration shape systematic is used to account for uncertainty in the relative calibration process (Sec. 3.2.1) particularly the difference in response close to the ends of cells. Linear fits to data/MC ratios of the PEcorr/cm vs. W in each view were used to define an uncertainty as a function of W, with different slopes in the middle of the detector vs. the edges [127].

## Muon Energy Scale

We account for additional uncertainty in the muon energy that comes from uncertainty in determining the muon track length. Separate systematics are considered which cover different correlated and uncorrelated effects, and are applied by varying the muon track

length in simulation, and redoing the reconstruction which results in a varied energy estimate. The first three cover effects that are uncorrelated between the FD, ND, and muon catcher, and have uncertainties of 0.13% (ND), 0.15% (FD), and 0.51% (muon catcher). These are dominated by detector mass accounting, primarily uncertainty in the exact volume/density of the FD scintillator and ND PVC mass [128]. The next is a correlated systematic that has uncertainties of 0.74% in the ND and FD, and 0.13% for the muon catcher<sup>5</sup>, which are dominated by muon range uncertainty coming from GEANT4, as well as mass accounting [128]. Finally, there is a systematic to deal with neutron pile-up at the ND, set as a one-sided +7 MeV up-scaling, and a systematic for uncertainty in the true initial energy of muons, which uses a one sided -9 MeV down-scaling [128].

### Geant4 Reweighting

New for this analysis, Geant4Reweighting [129] is an event reweighting framework that allows us to quantify the effects of changing our Geant4 hadron model without needing to re-run simulation. There are 11 systematic knobs covering uncertainties on inelastic scattering of hadrons produced in primary neutrino interactions as well as subsequent interactions of the particles in the detector [105]. They are incorporated into the analysis via a PCA similar to the GENIE and Flux systematics. The main difference is that the covariance matrices are constructed in reconstructed instead of true neutrino energy bins. This is because the systematic knobs only affect the particles produced in the interaction and do not affect the true energy of the neutrino. Three PCs are used, with a 100% scaling factor applied to ensure proper sensitivity coverage [130].

### Detector Aging

The detector aging effect, discussed in Sec. 3.14, is covered by a systematic sample which models the degradation of the scintillator as a simple linear decrease in the light model as a

---

<sup>5</sup>This means for an  $n\sigma$  shift up, each detector has their track lengths shifted up by  $n \times$  the corresponding percent.

function of time. The rate of 4.5% per year was chosen to match the decrease in NHit shifts seen in data, and is offset by a corresponding increase in the calibration constants which preserves muon response [111]. The effect can only decrease the scintillation efficiency, so this is implemented as a one-sided systematic, only allowing positive shifts. The effect is expected to be the same in both detectors so it is also a fully correlated between detectors.

#### 5.5.4 Other Systematics

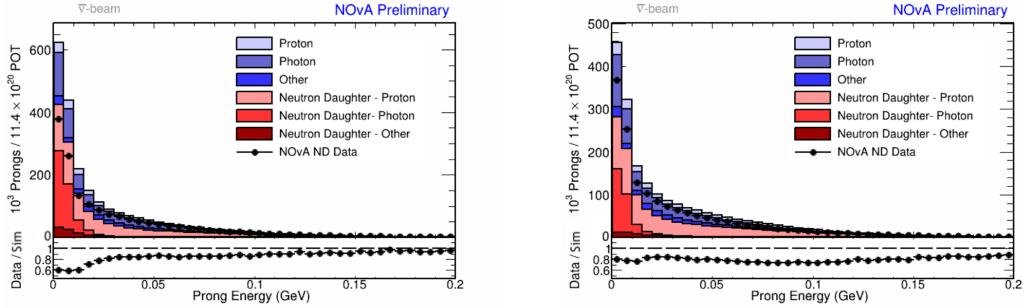
This section contains the remaining systematics that don't fit neatly into one of the other categories, or encompass several effects.

#### Neutron Model - MENATE

Neutrons are commonly produced in antineutrino interactions, so it is important we understand their uncertainty in our model. We account for neutron mismodeling using a systematic sample that replaces our baseline Geant4 simulation with one that uses a different neutron model, MENATE [131, 132], which includes additional data-derived cross sections for interactions on carbon. MENATE improves data/MC shape agreement in several neutron prong variables, including the prong energy as shown in Fig. 5.19. There are still some residual scaling differences, so an additional 33% scale factor is applied to the systematic to cover the difference [133]. Since the model can only be turned on or off, the systematic is included as a one-sided  $+1\sigma$  shift.

#### Cosmic Scale

The cosmic scale systematic is based on the statistical uncertainty of the cosmic data samples added into the predictions. For the  $\nu_e$  samples, the  $+/- 1\sigma$  error in each bin of the cosmic distribution is taken as the corresponding range for a Poisson distribution with mean equal to the unscaled cosmic events in that bin [111]. The error is then scaled down by the same amount as the events to match the beam livetime before being added in to the predictions. For the  $\nu_\mu$  samples, the shape of the energy distribution has more of an effect



**Figure 5.19:** Neutron prong energy for standard (left) and MENATE-supplemented Geant4 simulation (right). Prongs come from a neutron enhanced sub-sample selected from the ND RHC  $\bar{\nu}_\mu$  sample (the largest sample of neutrons available). Simulation is broken down by immediate particle type that produced the prong, with neutron related prongs in warm tones, and other primary particles in cool tones. From [133].

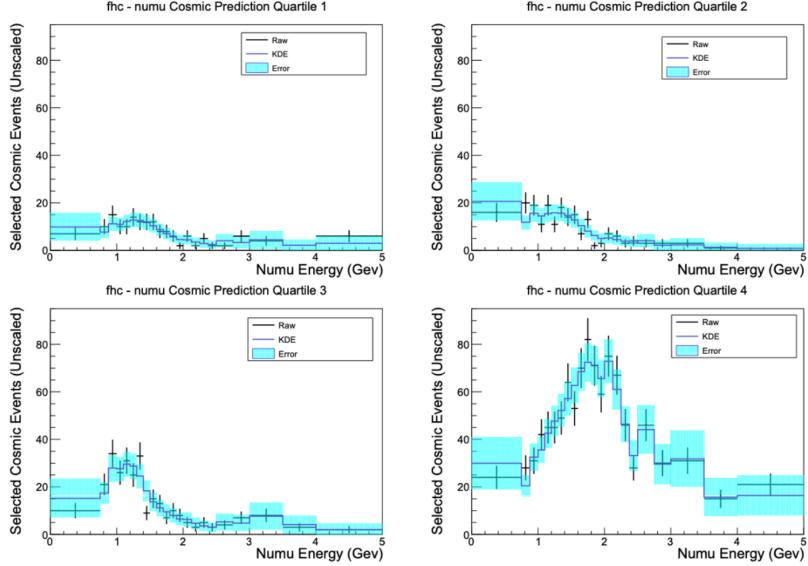
on results, so we smooth the cosmic spectra using a Kernel Density Estimate (KDE), which is a standard statistical technique for approximating the shape of unknown distributions [122]. An example of the smoothed distributions is shown in Fig. 5.20 for the FD FHC  $\nu_\mu$  quartiles.

#### 5.5.4.1 Normalization

The normalization systematic consists of several small systematic effects that affect the overall normalization of our final predictions and are not accounted for elsewhere. These are an uncertainty on the POT of 0.55%, an uncertainty in ND/FD mass differences of 0.19%, a simulation uncertainty of 0.4% meant to account for the impact of air bubbles in the liquid scintillator not simulated in GEANT4, and an uncertainty of 0.5% in the reduction in efficiency from pile-up in the ND [134]. The individual effects are added in quadrature to get a final normalization uncertainty of 0.86%.

#### 5.5.4.2 Michel Tagging

The efficiency with which we tag Michel electrons can affect the  $\nu_e$  background estimate through the Michel decomposition process. This is a small effect, so instead of creating



**Figure 5.20:** Selected cosmic events and their corresponding KDE for the FD FHC  $\nu_\mu$  quartiles, before scaling to match beam livetimes. The KDE error is shown in blue, and the Poisson error is shown as the black bars on the data points. From [105].

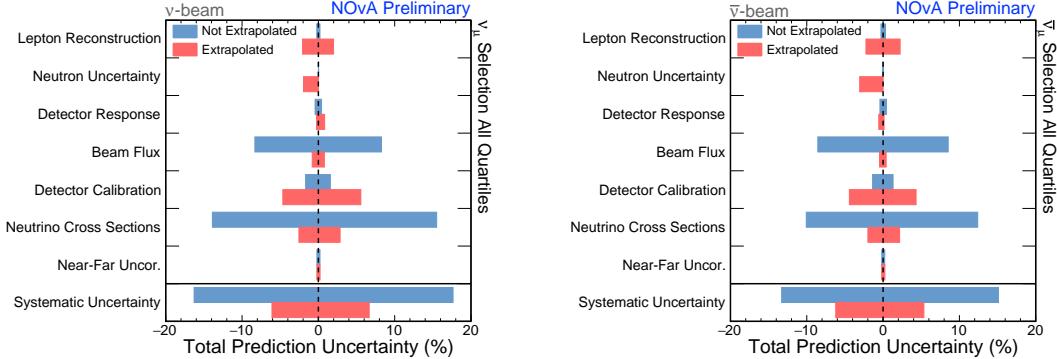
new systematic MC samples that vary the efficiency (a computationally costly endeavor), the tagging efficiency is artificially varied by randomly forcing 5% of Michels to pass/fail selection cuts for the  $\pm 1\sigma$  shifts. Carrying this through to the final  $\nu_e$  predictions results in a 1-2% shift in the  $\nu_e$  background [135].

### 5.5.5 Summary of uncertainties and the effect of extrapolation

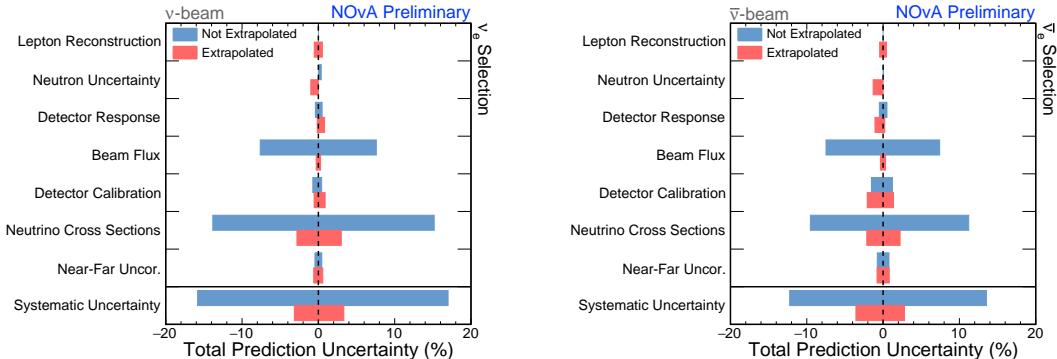
The effect of the systematic uncertainties on the final FD predicted event counts before and after extrapolation are summarized in Fig. 5.21 for the  $\nu_\mu$  sample, and Fig. 5.22 for the  $\nu_e$  sample, for both FHC and RHC. These illustrate the overall reduction in systematic uncertainty that extrapolation achieves, particularly for the large correlated systematics such as flux and cross-section. There are a few systematics for which extrapolation increases the uncertainty, such as lepton reconstruction <sup>6</sup>. This comes from introducing ND information with a higher degree of uncertainty into the FD predictions, but is more than

<sup>6</sup>This is referred to as ‘‘Muon Energy Scale’’ systematic in the section above. Additionally, ‘‘Near-Far Uncorr’’ refers to the normalization systematic.

made up for by the reduction in the largest uncertainties.



**Figure 5.21:** Summary of errors on the integrated number of  $\nu_\mu$  selected events for the FHC (left) and RHC (right) samples, with all quartiles summed. Errors bars are shown for predictions before (blue) and after (red) extrapolation is applied, and represent the  $\pm 1\sigma$  ranges for each systematic sample. The total error, taken as the sum in quadrature of the individual errors, is shown at the bottom. Predictions were generated using the best-fit oscillation parameters found in Sec. 7.2.



**Figure 5.22:** Summary of errors on the integrated number of  $\nu_e$  selected events for the FHC (left) and RHC (right) samples. Errors bars are shown for predictions before (blue) and after (red) extrapolation is applied, and represent the  $\pm 1\sigma$  ranges for each systematic sample. The total error, taken as the sum in quadrature of the individual errors, is shown at the bottom. Predictions were generated using the best-fit oscillation parameters found in Sec. 7.2.

## 5.6 Fitting Oscillation Parameters

Historically, NOvA has used a frequentist statistical approach to fitting oscillation parameters [113, 136, 137]. In 2023, a Bayesian approach to inference was used to analyze the NOvA data for the first time [138]. This analysis contains results that employ both approaches, which are described below. In both cases, the oscillation parameters are found by minimizing the binned Poisson log-likelihood function

$$-2 \ln L(\vec{\theta}, \vec{\phi}) = -2 \sum_{i=1}^N \left[ E_i(\vec{\theta}, \vec{\phi}) - O_i + O_i \ln \frac{O_i}{E_i(\vec{\theta}, \vec{\phi})} \right] + \sum_{j=1}^M \frac{\phi_j^2}{\sigma_j^2}, \quad (5.6)$$

where  $E_i(\vec{\theta}, \vec{\phi})$  are the predicted event count in analysis bin  $i$ , as a function of oscillation parameters  $\vec{\theta}$  and systematic uncertainties  $\vec{\phi}$ , and  $O_i$  are the observed data in analysis bin  $i$ . The second term is a sum is over the systematic uncertainties, computing a penalty term for each, where  $\phi_j$  and  $\sigma_j$  are the values and  $1\sigma$  ranges of systematic parameter  $j$ . These terms serve to keep the systematic parameters from being pulled too far away from their central values, with larger pulls incurring a larger likelihood penalty.

The simulated model used to predict the event counts at the FD uses the full oscillation probability in matter described in Sec. 1.2. The NOvA baseline  $L$  and density of the Earth  $\rho$  are held fixed at

$$L = 810 \text{ km}$$

$$\rho = 2.74 \text{ g/cm}^3 .$$

The solar oscillation parameters, which NOvA is not sensitive to, are held fixed at the values determined by the PDG [79]

$$\Delta m_{21}^2 = 7.53 \times 10^{-5} \text{ eV}^2$$

$$\sin^2(\theta_{12}) = 0.307 .$$

The remaining oscillation parameters,  $\Delta m_{32}^2$ ,  $\sin^2 \theta_{23}$ ,  $\delta_{\text{CP}}$ , and  $\sin^2 2\theta_{13}$  are allowed to vary in the fit, along with the systematic uncertainties.

Our sensitivity to  $\sin^2 2\theta_{13}$  is relatively low due to degenerate oscillation effects from the other parameters. At the same time, this parameter has been strongly constrained by the Daya Bay reactor neutrino experiment [139]

$$\sin^2 2\theta_{13} = 0.0851 \pm 0.0024.$$

This recent result is the most precise to date, and the data release for the publication included a 2D  $\chi^2$  surface for their measurement of  $\Delta m_{32}^2$  and  $\sin^2 2\theta_{13}$  as well, which makes it straightforward to include these as additional constraints in our fits. We consider several constraint options: no constraint (NOvA-only measurement), a 1D constraint on  $\sin^2 2\theta_{13}$ , or a 2D constraint on  $\sin^2 2\theta_{13}$  and  $\Delta m_{32}^2$ . The details of the implementation are specific to each statistical approach described below.

### 5.6.1 Frequentist Techniques

In the frequentist approach to fitting, we consider the log-likelihood as a  $\chi^2$  statistic

$$\chi^2(\vec{\theta}, \vec{\phi}) = -2L(\vec{\theta}, \vec{\phi}) \quad (5.7)$$

and the vector of parameters that minimizes this  $\chi^2$  is taken as the “best-fit” oscillation parameters. This minimum point, and its associated parameters, are referred to as the “best-fit” point  $\chi_{best}^2$ . Along with this best-fit point we present 1D confidence intervals, and 2D confidence regions, which are plots showing the regions of parameter space that, under many repeated independent trials of the experiment, we would expect to contain the true value of the parameters a certain percentage of the time. Typically these are drawn as contours indicating the  $1\sigma$ ,  $2\sigma$ , &  $3\sigma$  (68.27%, 95.45%, 99.73%) confidence levels. These regions are defined by means of the test statistic

$$\Delta\chi^2(\vec{\theta}, \vec{\phi}) = \chi^2(\vec{\theta}, \vec{\phi}) - \chi_{best}^2. \quad (5.8)$$

At each point in the phase space under consideration (for example, a 2D plot of  $\Delta m_{32}^2$  vs.  $\sin^2 \theta_{23}$ ) we calculate  $\Delta\chi^2$  using the parameter values at that point, while fitting the remaining oscillation and systematic parameters, to find the combination that gives the lowest  $\Delta\chi^2$ . This fit over the parameters not on the axes is referred to as “profiling” the parameters. In the Gaussian approximation, there are then well-defined cutoffs of  $\Delta\chi^2$  that correspond to  $1\sigma$ ,  $2\sigma$ , &  $3\sigma$  levels. For example, for a  $\Delta\chi^2$  with two degrees of freedom they are  $\Delta\chi^2 < 2.30$ ,  $\Delta\chi^2 < 6.18$ , and  $\Delta\chi^2 < 11.83$  respectively [23]. However, for several reasons, our experiment does not fall under the Gaussian approximation. Many of the bins, particularly near the disappearance maximum, can have very low event counts and are subject to Poisson fluctuations. Our parameter space also contains effective physical boundaries such as  $\sin^2 \theta_{23}$  restricted between 0 and 1, and a poorly-constrained cyclic parameter in  $\delta_{\text{CP}}$ . Using the Gaussian  $\Delta\chi^2$  cutoffs in this context will likely result in the confidence intervals and contours containing the true values of the parameters at a different rate than the stated  $1\sigma$ ,  $2\sigma$ , &  $3\sigma$  coverage [111].

To determine the cutoffs in  $\Delta\chi^2$  that do ensure proper  $1\sigma$ ,  $2\sigma$ , &  $3\sigma$  coverage we use the unified method of Feldman and Cousins [140]. At each point in the parameter space under consideration we create an ensemble ( $O(10^4)$ ) of mock experiments with statistical fluctuations applied to each bin. We calculate the  $\Delta\chi^2$  for each experiment and histogram them. We then find the  $\Delta\chi^2$  of the mock experiment corresponding to the quantile of the desired confidence level (e.g. 90%), and subtract from it the lowest  $\Delta\chi^2$  out of all the experiments (i.e. the overall best fit of all possibilities). This gives the  $\Delta\chi^2$  cutoff value for the desired confidence level for that point in parameter space. This process is repeated for all points in the space we are considering. For example, for a 2D plot of  $\sin^2 \theta_{23}$  vs.  $\delta_{\text{CP}}$  we may consider a  $30 \times 30$  grid of points in the expected range for each parameter. We can then draw the confidence level contours for the actual  $\Delta\chi^2$  using the “corrected” cutoff for each grid point on the plot. One other consideration is what to do about the nuisance parameters (i.e. the systematic and oscillation parameters not on the axes). To deal with these, we profile the parameters before throwing mock experiments [141].

Due to the large number of sample generations and fits required, this process is carried out on supercomputers at the National Energy Research Scientific Computing Center (NERSC), which brings the time to results down from months to weeks [142].

Finally, when it comes to the Daya Bay constraints, the 1D constraint on  $\sin^2 2\theta_{13}$  is implemented as a Gaussian penalty term in the fit. This gets added on to the  $\chi^2$  giving

$$\chi^{2\prime} = \chi^2 + \frac{(x - \mu)^2}{\sigma^2} , \quad (5.9)$$

where  $\mu$  is the result (0.0851),  $\sigma$  is the error on the result (0.0024), and  $x$  is the value of  $\sin^2 2\theta_{13}$  being considered. For the 2D Daya Bay constraint on  $\Delta m_{32}^2$  and  $\sin^2 2\theta_{13}$ , we simply add the  $\chi^2$  from the provided surface at the corresponding point in parameter space

$$\chi^{2\prime} = \chi^2 + \chi_{DB}^2 . \quad (5.10)$$

### 5.6.2 Bayesian Techniques

Bayesian statistics estimates parameters  $\vec{\theta}$  using Bayes theorem

$$P(\vec{\theta}|\vec{x}) = \frac{P(\vec{x}|\vec{\theta}) \times P(\vec{\theta})}{P(\vec{x})} , \quad (5.11)$$

where  $P(\vec{\theta}|\vec{x})$  is the probability that the parameters  $\vec{\theta}$  are true given the data  $\vec{x}$ , referred to as the posterior.  $P(\vec{x}|\vec{\theta})$  is the likelihood, which is found via Eq. 5.6:

$$\ln P(\vec{\theta}|\vec{x}) = -2 \ln L(\vec{\theta}, \vec{\phi}) . \quad (5.12)$$

$P(\vec{\theta})$  is the prior, which represents our knowledge of the parameters, and constrains the posterior accordingly. It is also possible to choose an “uninformed” prior that is constant across the range of the parameter. Finally,  $P(\vec{x})$  is the evidence, or probability of observing the data. It is independent of the parameters, and acts only as a normalization constant. Since the NOvA data is fixed, and our goal is to find the parameters that maximize

the posterior (not the absolute posterior probability), we can then ignore this term when estimating parameters [138]

$$P(\vec{\theta}|\vec{x}) \sim P(\vec{x}|\vec{\theta}) \times P(\vec{\theta}). \quad (5.13)$$

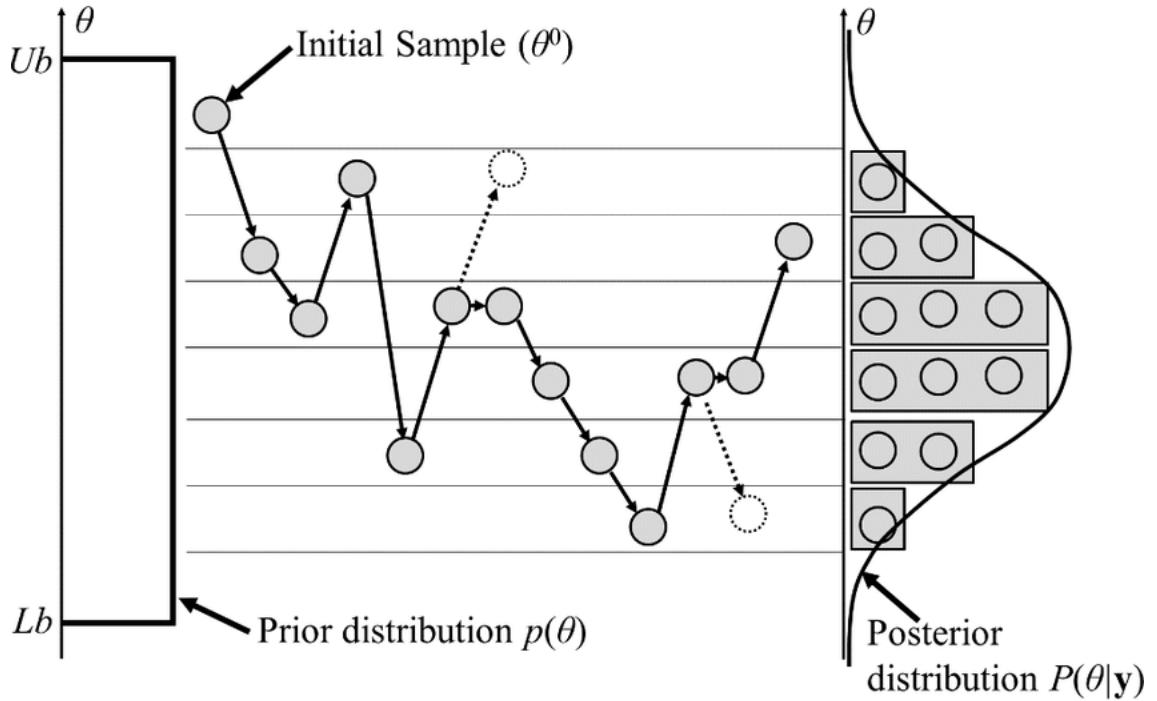
In Bayesian statistics we also wish to make statements about our degree of confidence in the values of the parameters that we find that maximize the posterior. We make 1D or 2D plots of the posterior probability for some chosen parameters, and define 1D “credible intervals”, or 2D “credible regions” which indicate the  $1\sigma$ ,  $2\sigma$ , &  $3\sigma$  regions containing the highest probability bins that yield 68.27%, 95.45%, & 99.73% of the total probability distribution respectively [143]. These credible regions differ from the frequentist confidence intervals in that they represent a region, under our specified model, that should contain the true value with the given probability [144].

When making these plots we also need to deal with the nuisance parameters. In frequentist statistics we profiled over the parameters not being considered, but in Bayesian statistics we instead “marginalize” over them by integrating the likelihood across them:

$$P(\vec{\theta}|\vec{x}) = \int P(\vec{\theta}, \vec{\delta} | \vec{x}) \, d\vec{\delta}. \quad (5.14)$$

Here  $\vec{\theta}$  are the parameters of interest and  $\vec{\delta}$  are the nuisance parameters.

The exact calculation of the posterior can quickly become difficult when the parameter space extends beyond just a few parameters [143]. With NOvA’s large number of systematics and oscillation parameters, a different analytical approach is necessary to estimate the shape of the posterior. We use Markov Chain Monte Carlo (MCMC) to solve this problem. The Markov Chain builds up a series of sampled points (the “chain”) where the next point is chosen from a probability distribution based on the current point. An example is shown in Fig. 5.23. This is an efficient way of estimating the posterior, and can achieve any desired degree of precision given enough samples [145]. NOvA uses the Metropolis-Rosenbluth-Rosenbluth-Teller-Teller-Hastings ( $MR^2T^2H$ ) sampling method [146], referred to as ARIA (for Arianna Rosenbluth, a major contributor).



**Figure 5.23:** An illustration of the MCMC sampling technique used by NOvA, for a 1D example scenario. The chain of points is built up by iteratively proposing a new point, sampled from a Gaussian distribution, and randomly deciding whether to accept it and add to the chain, or stay at the current point. The acceptance criteria prefers higher density regions based on the posterior, so with enough samples we build up the correct distribution shown on the right. From [147].

The Daya Bay constraints for Bayesian results are applied in a straightforward way by simply multiplying the posterior (with appropriate normalizations) by the likelihood of the external constraint [143]. So the MCMC sampling only needs to be done once to get a posterior without any constraint, and then we can get 1D or 2D constrained versions by simply multiplying by the likelihood of the corresponding Daya Bay constraint.

## Chapter 6

# Gaining Sensitivity

NOvA is now a mature experiment, with many years of data and a robust analysis procedure. Other than the inclusion of new data, we primarily expect gains in sensitivity to come from squeezing out any additional sensitivity that we can from the existing data. This means smaller analysis optimizations or additions, as opposed to large overhauls of the procedure. In the years leading up to this analysis we conducted several studies that examined the state of the 3-flavor analysis, and looked for additional ways in which we could improve our sensitivity to oscillation parameters. This chapter covers the results of these studies, some of which led to new additions in this analysis, while some could be targeted for future analyses. First we describe the implementation of the Daya Bay constraint on  $\sin^2 2\theta_{13}$  and  $\Delta m_{32}^2$ , and look ahead to how our sensitivity to the oscillation parameters might evolve over time. Then, we discuss efforts to reclaim both  $\nu_\mu$  and  $\nu_e$  events that were failing cuts, and the motivation for the new low-energy sample. Finally, we go over the many improvements that I made to the Michel electron reconstruction, and conclude with ideas for its future application to the 3-flavor analysis.

Since many of the studies in this section were conducted in years prior, before the details of this analysis had been finalized, they may use the datasets, selections, and/or systematics from NOvA’s 2020 analysis [113]. The overarching procedure is largely the same, but with some of the re-optimizations for the new simulation and other updates

detailed in previous chapters. A full summary of the 2020 analysis methodology can be found at [111].

## 6.1 Daya Bay Constraints and Future Sensitivity

One way we can improve our sensitivity to oscillation parameters is by incorporating results from other experiments as additional constraints in our fits. In past NOvA 3-flavor oscillation analyses, we used the PDG value of  $\sin^2 2\theta_{13}$  as a constraint, implemented as a Gaussian penalty term in the same way as the 1D constraint described in Sec. 5.6.1 [113]. However, in 2022 Daya Bay was expected to release a new precision measurement of  $\sin^2 2\theta_{13}$  alongside a measurement of the mass splitting  $\Delta m_{32}^2$ . We had previously only incorporated a constraint on  $\sin^2 2\theta_{13}$  and were interested in whether including a constraint on  $\Delta m_{32}^2$  as well could further improve our sensitivity. To gauge the impact of such a constraint, and understand its potential effect on future measurements, we compared fitting results with several different constraints at current and future POT levels.

### 6.1.1 Implementing the constraints

We considered several versions of a Daya Bay constraint, using their 2018 results [148] as the 2022 results hadn't been published yet. First is the nominal 1D PDG constraint on  $\sin^2 2\theta_{13}$  that NOvA had been using. Then, the two independent 1-Dimensional constraints using the Daya Bay results is implemented in the same way as the existing  $\sin^2 2\theta_{13}$  constraint. So, one constraint using their measurement for  $\sin^2 2\theta_{13}$ ,

$$\sin^2 2\theta_{13} = 0.0851 \pm 0.0024 \tag{6.1}$$

and a separate constraint using their measurements for  $\Delta m_{32}^2$  in each mass ordering

$$\Delta m_{32}^2 = 2.471 \pm 0.07 \times 10^{-3} \text{ eV}^2 \text{ (Normal Mass Ordering)} \tag{6.2}$$

$$\Delta m_{32}^2 = -2.575 \pm 0.07 \times 10^{-3} \text{ eV}^2 \text{ (Inverted Mass Ordering)} \quad (6.3)$$

[148]. Lastly, we applied a constraint that used the provided 2D  $\chi^2$  surface for their results. This required some additional treatment. Daya Bay does not measure  $\Delta m_{32}^2$  directly. At their short 1 km baseline, they instead measure an effective mass-squared difference  $\Delta m_{ee}^2$  related to the wavelength of the oscillation [139]. They measure the  $\bar{\nu}_e$  disappearance oscillation channel, which has an oscillation probability given by

$$\begin{aligned} P(\bar{\nu}_e \rightarrow \bar{\nu}_e) = & 1 - \cos^4 \theta_{13} \sin^2 2\theta_{12} \sin^2 \Delta_{21} \\ & - \sin^2 2\theta_{13} (\cos^2 \theta_{12} \sin^2 \Delta_{31} + \sin^2 \theta_{12} \sin^2 \Delta_{32}) , \end{aligned} \quad (6.4)$$

where  $\Delta_{ij} = 1.27 \Delta m_{ij}^2 L/E$  with  $L$  as the baseline and  $E$  as the neutrino energy. At their baseline, the two oscillation phases  $\Delta_{31}$  and  $\Delta_{32}$  are indistinguishable [149], and so the terms in the parentheses are approximated by the effective phase  $\Delta m_{ee}^2$ :

$$P(\bar{\nu}_e \rightarrow \bar{\nu}_e) = 1 - \cos^4 \theta_{13} \sin^2 2\theta_{12} \sin^2 \Delta_{21} - \sin^2 2\theta_{13} \Delta m_{ee}^2 . \quad (6.5)$$

Their measurement of  $\Delta m_{ee}^2$  can then be converted to  $\Delta m_{32}^2$  via the equation

$$|\Delta m_{32}^2| = |\Delta m_{ee}^2| - \alpha \cos^2(\theta_{12}) \Delta m_{21}^2 , \quad (6.6)$$

where  $\alpha$  is  $\pm 1$  in the normal/inverted mass ordering, and the solar oscillation parameters are fixed at their PDG values. Their data is provided as a grid of points in  $\Delta m_{ee}^2$  vs.  $\sin^2 2\theta_{13}$  space, with each point giving their measured  $\chi^2$  value for that combination of parameters. To turn this into a usable constraint in our fits, we first convert the grid of points to a surface, using ROOT's TGraph2D class. This allows us to interpolate between points using the functions provided by the class, to get  $\chi^2$  values at any combination of parameters, a requirement in our fits. In the fits, for a given value of  $\Delta m_{32}^2$  and  $\sin^2 2\theta_{13}$ , we would first convert to  $\Delta m_{ee}^2$  using Eq. 6.6<sup>1</sup>, and then sample that point from the

---

<sup>1</sup>In the data release for their latest 2023 publication [139] they provided a grid of points in  $\Delta m_{32}^2$  vs.  $\sin^2 2\theta_{13}$  space as well, so we are able to skip the conversion step and just directly use that surface for the

TGraph2D object to get the correct  $\chi^2$ .

### 6.1.2 Impact on oscillation parameters

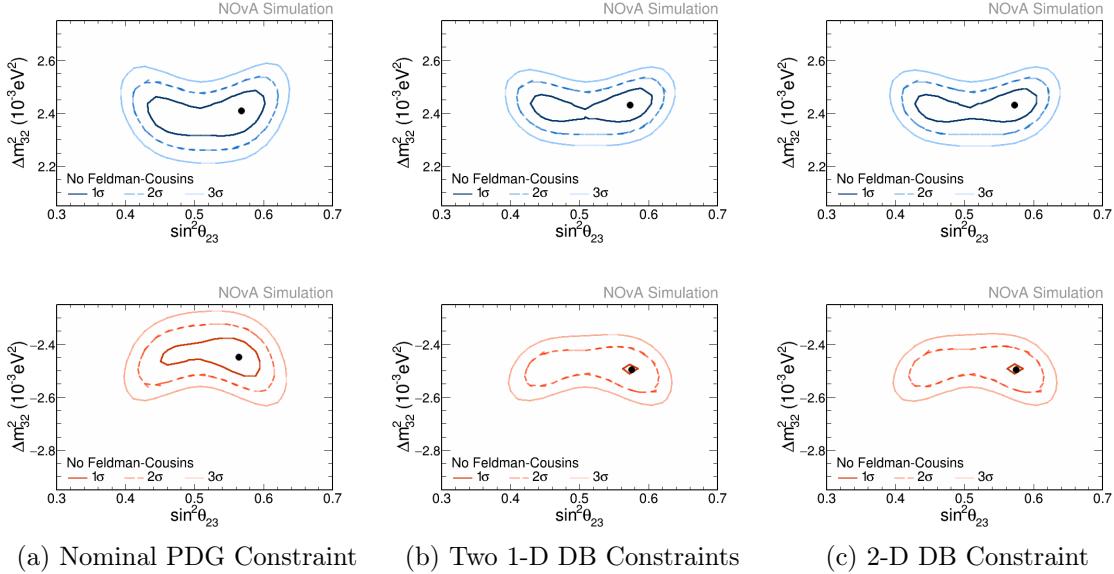
To gauge the impact of the constraints on our oscillation measurements, we construct separate 2D frequentist confidence regions using each constraint, and compare them. The FD predictions used in the fits here were made using the selections, binnings, and systematics from NOvA’s 2020 analysis [113]. When doing preliminary tests like this, we don’t use real data in order to avoid bias. Instead we generate “fake data” which is just a predicted FD spectrum at a particular value of the oscillation parameters that we use in place of real data in our fits. All fake data used in this section was generated at NOvA’s 2020 analysis best-fit point. Additionally, we do not apply Feldman-Cousins corrections at this testing stage due to their prohibitive computational cost. The contours shown here are not representative of final results, but meant to give an idea of the relative sensitivity to different analysis approaches.

Confidence regions in  $\Delta m_{32}^2$  vs.  $\sin^2 \theta_{23}$  space are shown in Fig. 6.1 using each version of the constraint. Both versions of the Daya Bay constraint improve the sensitivity compared to the nominal constraint, and by similar amounts. In both mass orderings, the contours are “squeezed” in the  $\Delta m_{32}^2$  direction, with the inverted ordering showing additional reduction in the  $\sin^2 \theta_{23}$  space, particularly at the  $1\sigma$  level. Fig. 6.2 shows the corresponding confidence regions in  $\sin^2 \theta_{23}$  vs.  $\delta_{\text{CP}}$  space. Again we see similar improvement from both versions of the Daya Bay constraint, with a more restricted contour in the inverted mass ordering. For both sets of plots, this enhanced restriction in the inverted mass ordering reflects the fact that both Daya Bay and NOvA prefer the normal mass ordering.

We also wanted to understand how much of the change we were seeing in the contours was coming specifically from the difference between NOvA and Daya Bay’s measured values of  $\Delta m_{32}^2$ , versus simply the inclusion of this additional data point in the fits. To test this, we imagined a scenario where Daya Bay would make a measurement of  $\Delta m_{ee}^2$  that agrees

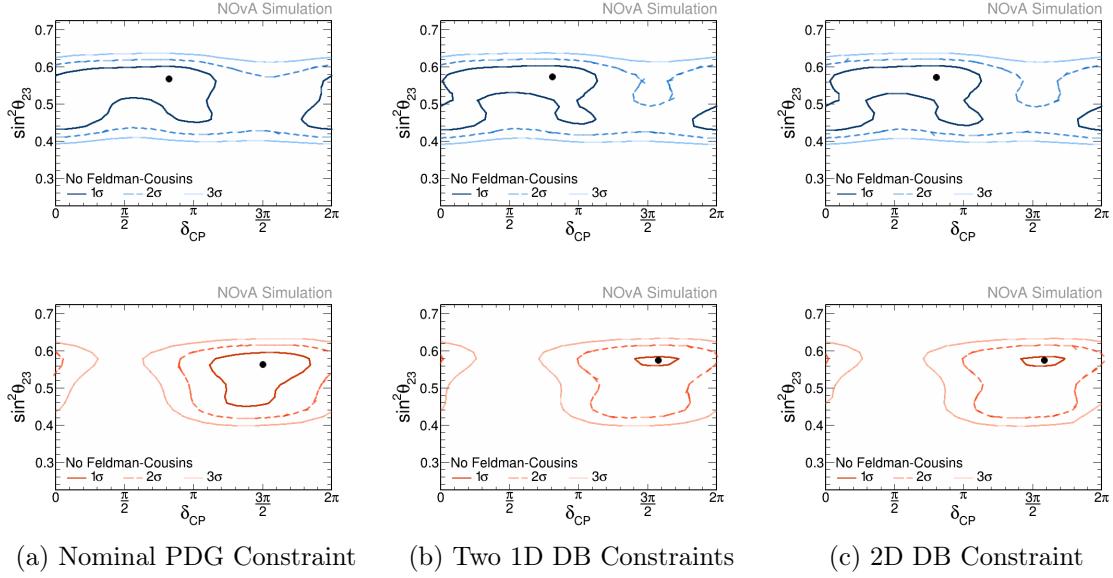
---

fit results later in this thesis (Chapter 7).



**Figure 6.1:** Fake data fits showing frequentist confidence-level contours for  $\Delta m_{32}^2$  vs.  $\sin^2 \theta_{23}$  in the normal (top) and inverted (bottom) mass orderings, using several different constraints on the value of  $\sin^2 2\theta_{13}$ . The left column uses the nominal 1D constraint on  $\sin^2 2\theta_{13}$  using the PDG value [150]. The middle column uses two uncorrelated 1D constraints on  $\sin^2 2\theta_{13}$  and  $\Delta m_{32}^2$ , and the right column uses a 2D correlated constraint on  $\Delta m_{32}^2$  and  $\sin^2 2\theta_{13}$ , all based on the 2018 Daya Bay measurement [148]. Black dots indicate the best-fit point in each plot.

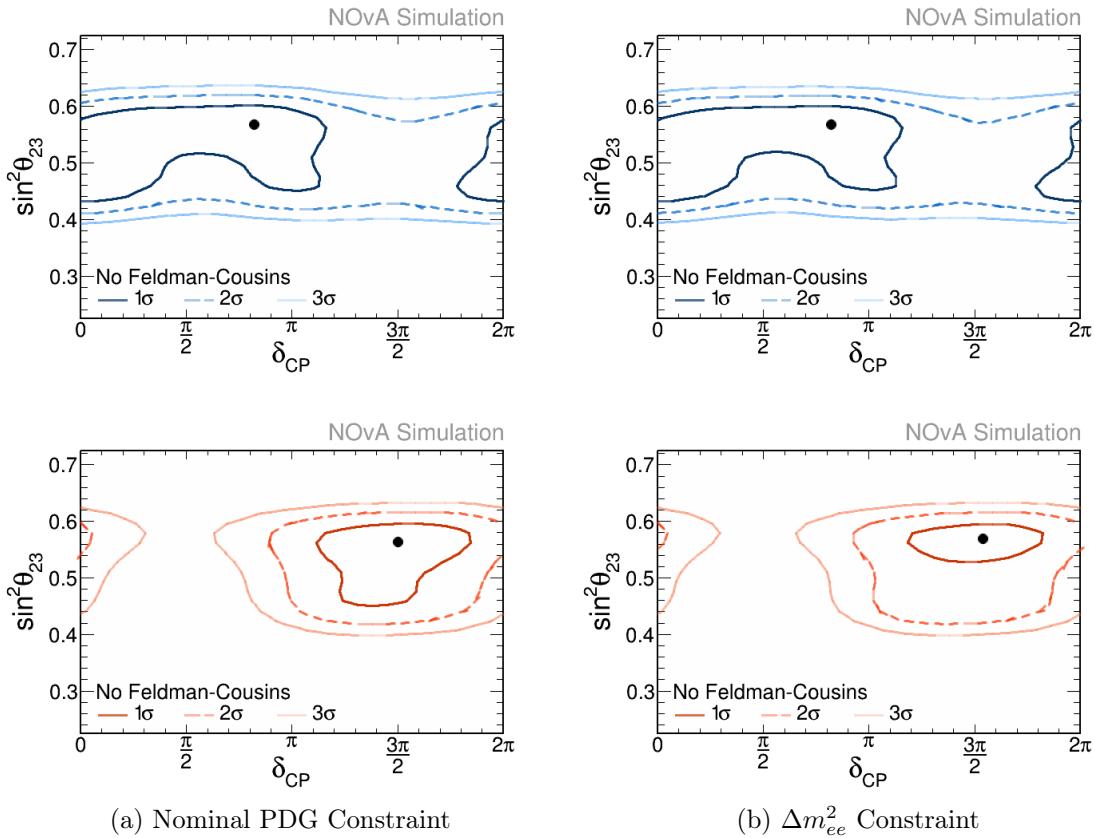
exactly with NOvA in the normal mass ordering, and differs in the inverted mass ordering. Then, we can compare improvements in each mass ordering to see how much the difference in the measurements affects the contours. Using Eq. 6.6, this would correspond to  $\Delta m_{ee}^2 = 2.46 \pm 0.07 \times 10^{-3} \text{ eV}^2$  which would give  $\Delta m_{32}^2 = 2.41 \pm 0.07 \times 10^{-3} \text{ eV}^2$  in the normal mass ordering (matching NOvA's 2020 result), and  $\Delta m_{32}^2 = -2.51 \pm 0.07 \times 10^{-3} \text{ eV}^2$  in the inverted mass ordering. We implemented these as 1D constraints on  $\Delta m_{32}^2$  along with the same 1D Daya Bay constraint on  $\sin^2 2\theta_{13}$  as before, and did fits to produce 2D confidence level contours to compare with the nominal case. Fig. 6.3 shows the  $\sin^2 \theta_{23}$  vs.  $\delta_{\text{CP}}$  contours and Fig. 6.4 shows the  $\Delta m_{32}^2$  vs.  $\sin^2 \theta_{23}$  contours. What we see is that in the normal mass ordering, the  $\Delta m_{ee}^2$  constraint gives less improvement over the nominal constraint, compared to Figs. 6.3, and 6.4. At the same time, the inverted mass ordering still experiences a larger reduction in the size of the contours (although not as



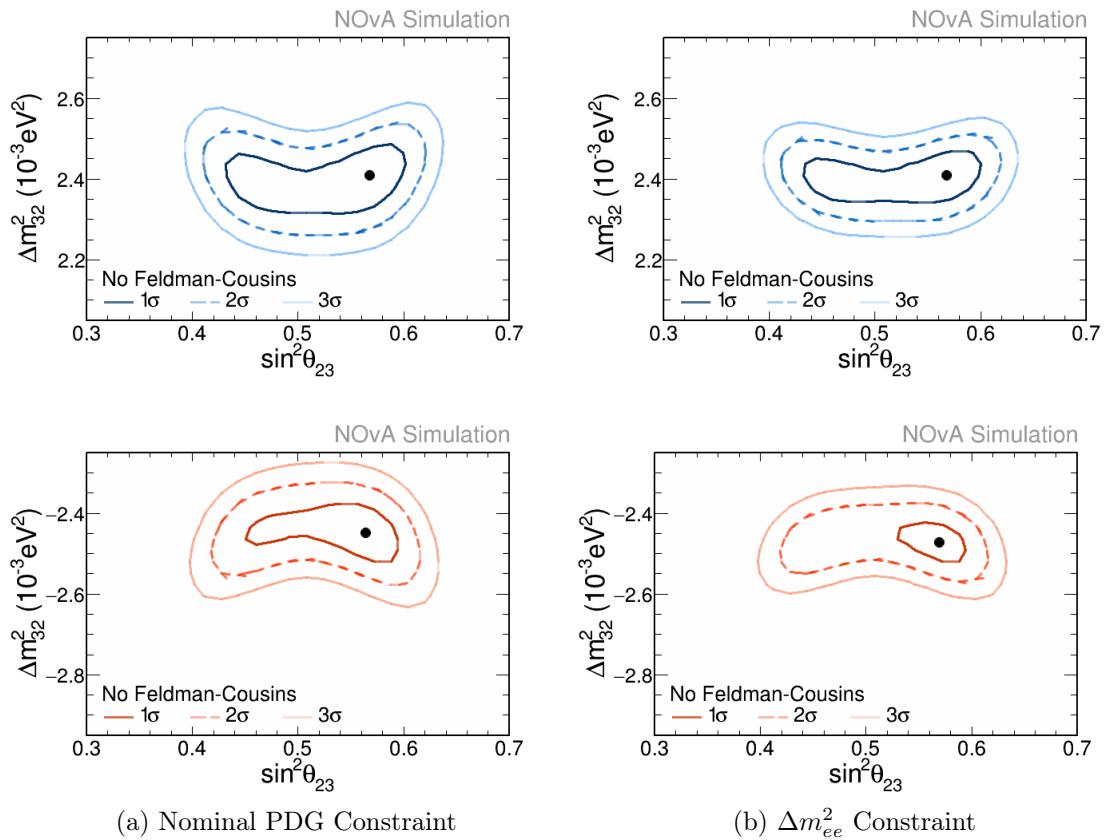
**Figure 6.2:** Fake data fits showing frequentist confidence-level contours for  $\sin^2 \theta_{23}$  vs.  $\delta_{CP}$  in the normal (top) and inverted (bottom) mass orderings, using several different constraints on the value of  $\sin^2 2\theta_{13}$ . The left column uses the nominal 1D constraint on  $\sin^2 2\theta_{13}$  using the PDG value [150]. The middle column uses two uncorrelated 1D constraints on  $\sin^2 2\theta_{13}$  and  $\Delta m_{32}^2$ , and the right column uses a 2D correlated constraint on  $\Delta m_{32}^2$  and  $\sin^2 2\theta_{13}$ , all based on the 2018 Daya Bay measurement [148]. Black dots indicate the best-fit point in each plot.

large as in Figs. 6.3, and 6.4). This suggests that the tension between the NOvA and Daya Bay results does play a role in constraining the shape of the contours along with overall sensitivity improvements from including additional information in the fit.

Given the potential improvement seen from using the additional Daya Bay constraints on  $\Delta m_{32}^2$ , we elected to include the Daya Bay 2D constraint as an additional option in our fits going forward. The 2D constraint was selected over the two 1D constraint since they gave very similar results and the 2D version better represents the actual Daya Bay measurement, and would preserve any small correlations between the parameters. The effect of its inclusion in the final results for this analysis is shown in Chapter 7.



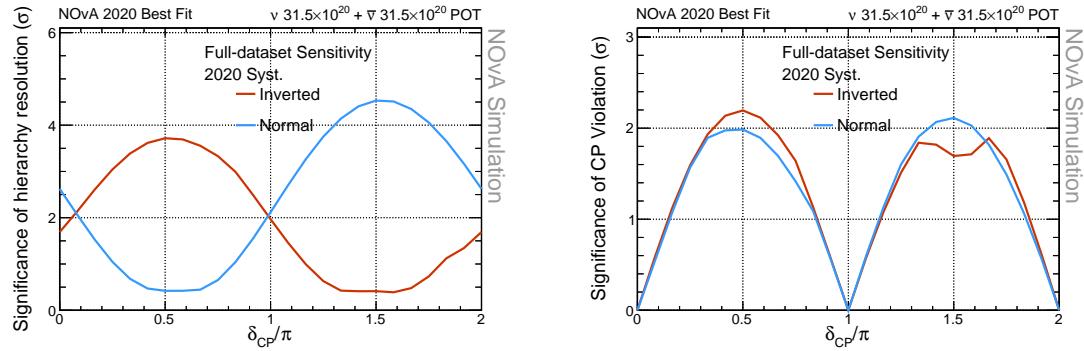
**Figure 6.3:** Fake data fits showing frequentist confidence-level contours for  $\sin^2 \theta_{23}$  vs.  $\delta_{CP}$  in the normal (top) and inverted (bottom) mass orderings, using different constraints. The left column uses the nominal 1D constraint on  $\sin^2 2\theta_{13}$  using the PDG value [150]. The right column implements two uncorrelated 1D constraints, using Daya Bay's result [148] for  $\sin^2 2\theta_{13}$ , and a  $\Delta m_{ee}^2$  constraint which agrees with NOvA's  $\Delta m_{32}^2$  value in the normal ordering and differs in the inverted ordering. Black dots indicate the best-fit point in each plot.



**Figure 6.4:** Fake data fits showing frequentist confidence-level contours for  $\Delta m_{32}^2$  vs.  $\sin^2 \theta_{23}$  in the normal (top) and inverted (bottom) mass orderings, using different constraints. The left column uses the nominal 1D constraint on  $\sin^2 2\theta_{13}$  using the PDG value [150]. The right column implements two uncorrelated 1D constraints, using Daya Bay's result [148] for  $\sin^2 2\theta_{13}$ , and a  $\Delta m_{ee}^2$  constraint which agrees with NOvA's  $\Delta m_{32}^2$  value in the normal ordering and differs in the inverted ordering. Black dots indicate the best-fit point in each plot.

### 6.1.3 Future Sensitivities

We were also interested in how much our sensitivity to different measurements would improve with more data, and whether we would become systematics-limited at any point. At the time it was expected that NOvA would take around  $60 \times 10^{20}$  POT of data across the lifetime of the experiment, split equally between FHC and RHC. A study was conducted to quantify our sensitivity at these final POT levels. Fig. 6.5 shows the projected sensitivity to reject the incorrect mass ordering and measure CP violation, while Fig. 6.6 shows the sensitivity to reject maximal mixing, and measure the octant of  $\theta_{23}$ . Both figures are made assuming a final NOvA dataset of  $63 \times 10^{20}$  POT, split evenly between FHC and RHC.

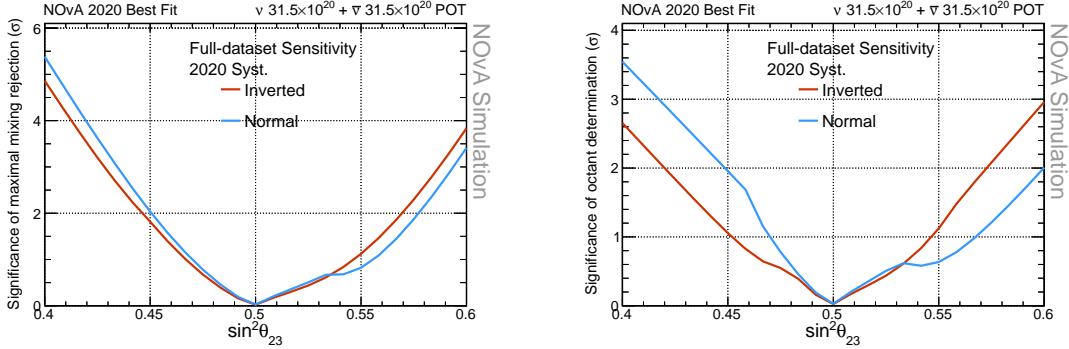


**Figure 6.5:** The significance at which we could reject the incorrect mass hierarchy (left) or measure CP violation (right) as a function of the true value of  $\delta_{CP}$ , for a future  $63 \times 10^{20}$  POT level, with data split evenly between FHC and RHC. From [151].

At each point on each of the curves, fake data is generated setting the variable in question ( $\delta_{CP}$  or  $\sin^2 \theta_{23}$ ) to the value on the x-axis, setting the sign of  $\Delta m_{32}^2$  to match the current mass-ordering, with all remaining parameters set to match the NOvA 2020 best-fit values.

We then perform two fits to this fake data, a restricted and unrestricted fit, and quote the square root of the difference in their  $\chi^2$  as the sensitivity. The unrestricted fit allows all the fit parameters to vary freely, and acts as the “best-fit” for each point. The restricted fit is done keeping the range of one of the parameters fixed to the “wrong” region of phase space.

$$\sigma = \sqrt{\chi^2_{restricted} - \chi^2_{best-fit}} \quad (6.7)$$

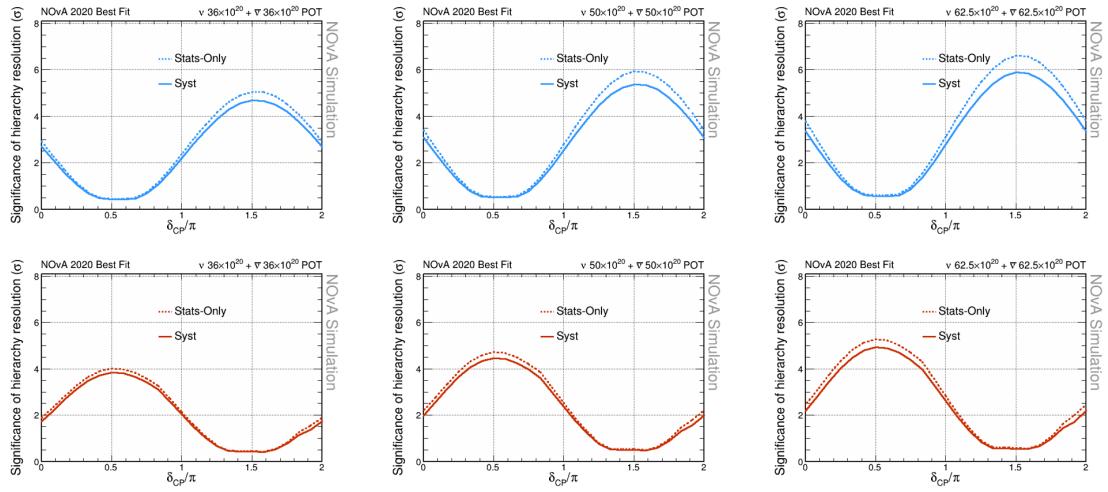


**Figure 6.6:** The significance at which we could reject maximal mixing (left) and determine the octant of  $\theta_{23}$  (right) as a function of the true value of  $\theta_{23}$ , for a future  $63 \times 10^{20}$  POT level, with data split evenly between FHC and RHC. From [151].

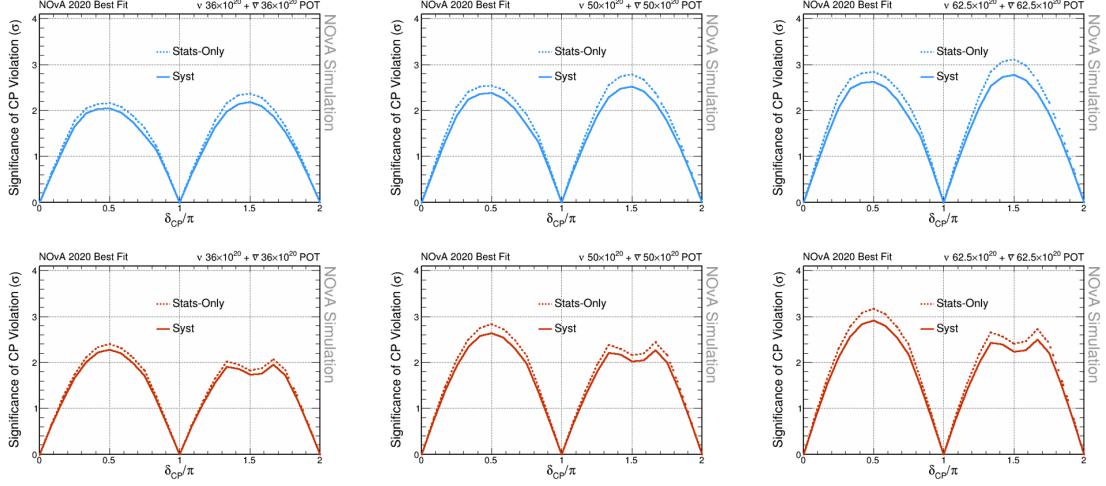
For the mass-ordering sensitivity the restricted fit only allows values of  $\Delta m_{32}^2$  in the opposite mass-ordering than what was used to generate the fake data. So for the inverted mass ordering curve, the fits are restricted to the normal mass ordering, and vice versa. The CP violation sensitivity restricts  $\delta_{\text{CP}}$  to be either 0 or  $\pi$  (i.e. no CP violation). For the maximal mixing rejection sensitivity,  $\sin^2 \theta_{23}$  is restricted to a small region around 0.5, the maximal mixing value. The octant determination sensitivity restricts  $\theta_{23}$  to the opposite octant from the current point used to generate fake data. As we approach the restricted regions on the x-axis, our “best-fit”  $\chi^2$  at each point will naturally agree with the restricted fits, and our rejection sensitivity goes to zero. This makes sense, as we cannot reject maximal mixing if its true, and we cannot measure any CP violation if there is none. The one exception is the mass-ordering sensitivity which doesn’t restrict the value of  $\Delta m_{32}^2$ , only the sign. Here we see a maximum sensitivity in the regions of  $\delta_{\text{CP}}$  that give the highest  $\nu_e$  ( $\bar{\nu}_e$ ) asymmetry for a given mass ordering.

We then extended this study out to higher POT levels, in order to see whether the gains in sensitivity leveled off at any point, indicating that our measurements had become systematically-limited instead of statistically-limited. We considered 72, 100, and 125  $\times 10^{20}$  POT, the largest amount being double the expected full NOvA dataset at the time. Figs. 6.8, 6.7, 6.9, and 6.10 show the sensitivities. In each plot, the dotted lines indicate fits

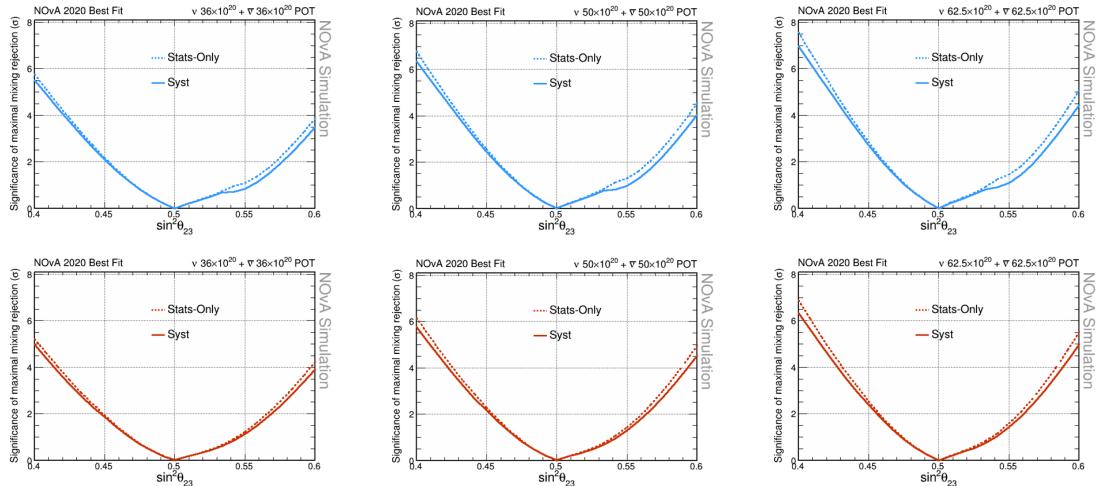
without including the effects of systematic uncertainties (so-called “stats-only” plots), while the fits with the full systematic treatment are shown in solid lines. The sensitivity is lower when accounting for systematic uncertainties but we still continue to see improvement with additional POT in all cases. This is an indicator that we are unlikely to become systematics-limited given our current uncertainty levels.



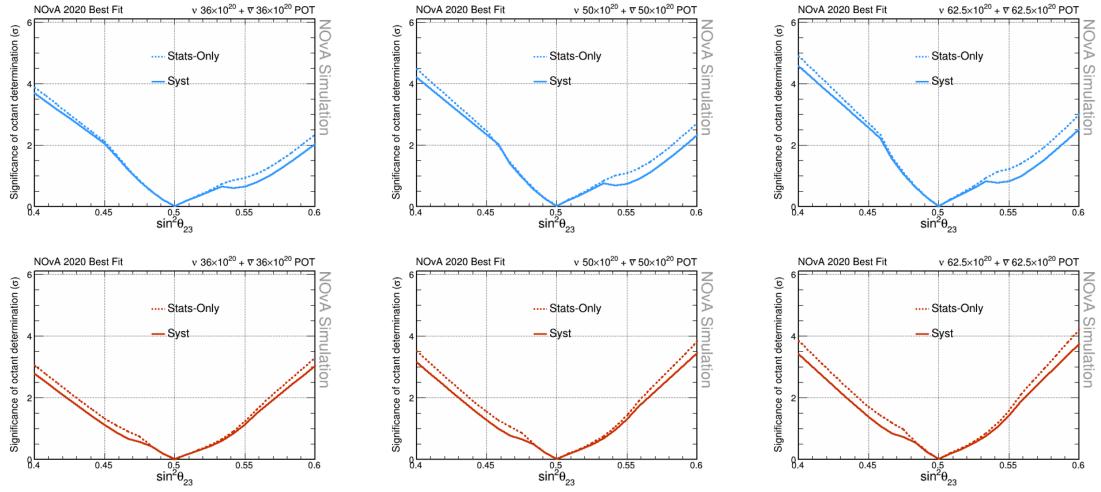
**Figure 6.7:** NOvA’s sensitivity to determine the neutrino mass ordering at increasing levels of POT, for the normal (top) and inverted (bottom) mass orderings, as a function of the true value of  $\delta_{CP}$ . 72, 100, and  $125 \times 10^{20}$  POT scenarios are shown, with data split evenly between FHC and RHC. Dotted lines indicate fits without systematics included, while solid lines have all systematics included.



**Figure 6.8:** NOvA’s sensitivity to discover CP violation at increasing levels of POT, for the normal (top) and inverted (bottom) mass orderings, as a function of the true value of  $\delta_{CP}$ .  $72, 100$ , and  $125 \times 10^{20}$  POT scenarios are shown, with data split evenly between FHC and RHC. Dotted lines indicate fits without systematics included, while solid lines have all systematics included.

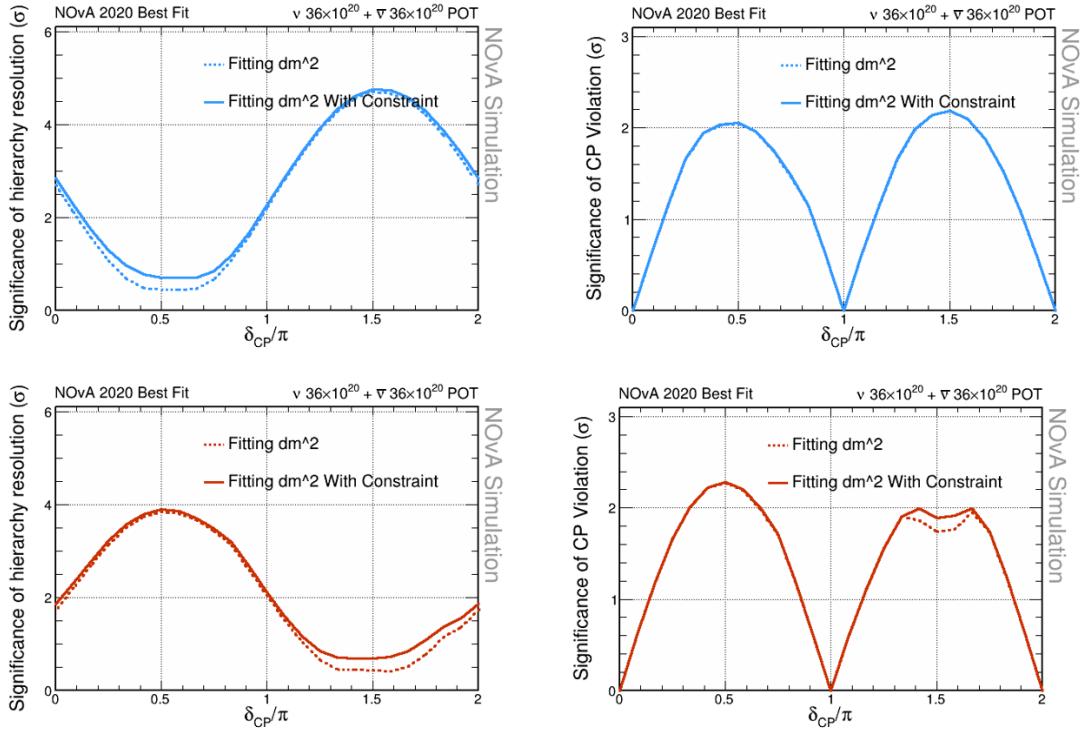


**Figure 6.9:** NOvA’s sensitivity to reject maximal mixing at increasing levels of POT, for the normal (top) and inverted (bottom) mass orderings, as a function of the true value of  $\delta_{CP}$ .  $72, 100$ , and  $125 \times 10^{20}$  POT scenarios are shown, with data split evenly between FHC and RHC. Dotted lines indicate fits without systematics included, while solid lines have all systematics included.

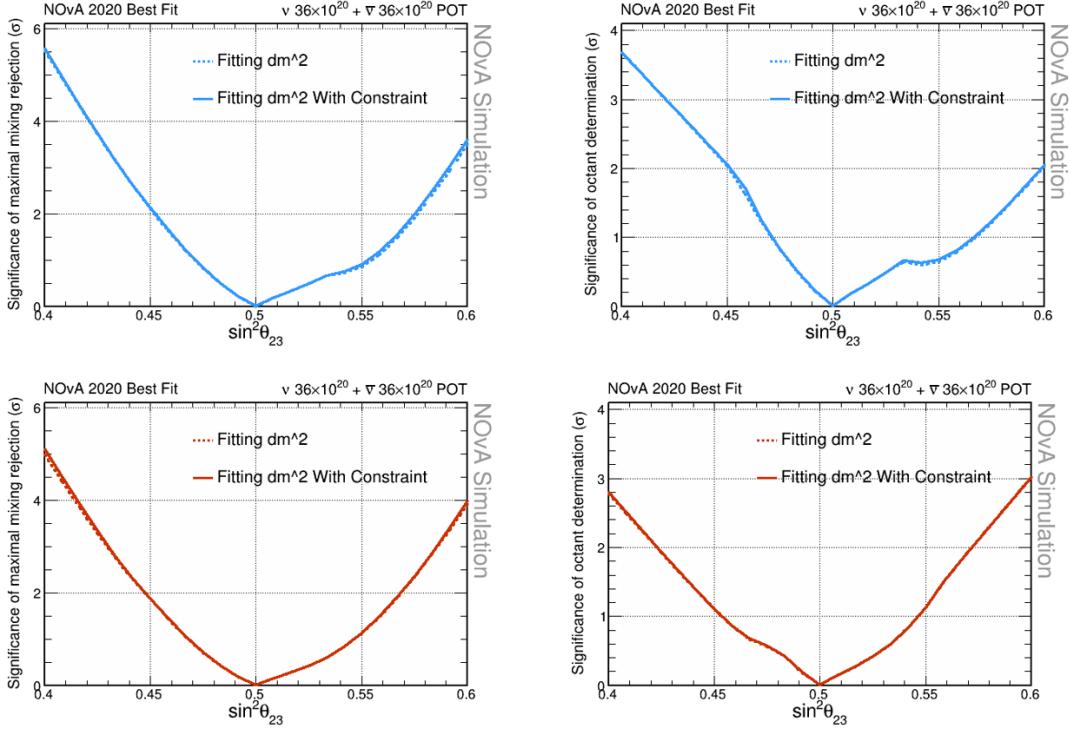


**Figure 6.10:** NOvA’s sensitivity to determine the octant of  $\theta_{23}$  at increasing levels of POT, for the normal (top) and inverted (bottom) mass orderings, as a function of the true value of  $\delta_{CP}$ .  $72, 100$ , and  $125 \times 10^{20}$  POT scenarios are shown, with data split evenly between FHC and RHC. Dotted lines indicate fits without systematics included, while solid lines have all systematics included.

Finally, we explored a future sensitivity scenario with an additional constraint on  $\Delta m_{32}^2$ . This was done to check whether including a constraint at a fixed uncertainty level might eventually negatively impact our results with further datataking. we considered a 1D gaussian constraint on  $\Delta m_{32}^2$  around the NOvA 2020 best-fit value with a 3% uncertainty (matching the uncertainty in the Daya Bay 2018 result), and reproduced the  $72 \times 10^{20}$  POT sensitivity plots with and without the constraint, shown in Figs. 6.11, 6.12. In all cases this additional constraint did not significantly impact the results, providing a mild improvement at a few values of  $\delta_{CP}$ .



**Figure 6.11:** NOvA's sensitivity to determine the mass ordering (left) and measure CP violation (right) for the normal (top) and inverted (bottom) mass orderings, as a function of the true value of  $\delta_{CP}$ . Plots shown for a future  $72 \times 10^{20}$  POT scenario, split evenly between FHC and RHC data. The dotted lines indicate the nominal fake data fits while the solid lines include an additional 3% constraint on  $\Delta m_{32}^2$  around NOvA's 2020 best-fit value.

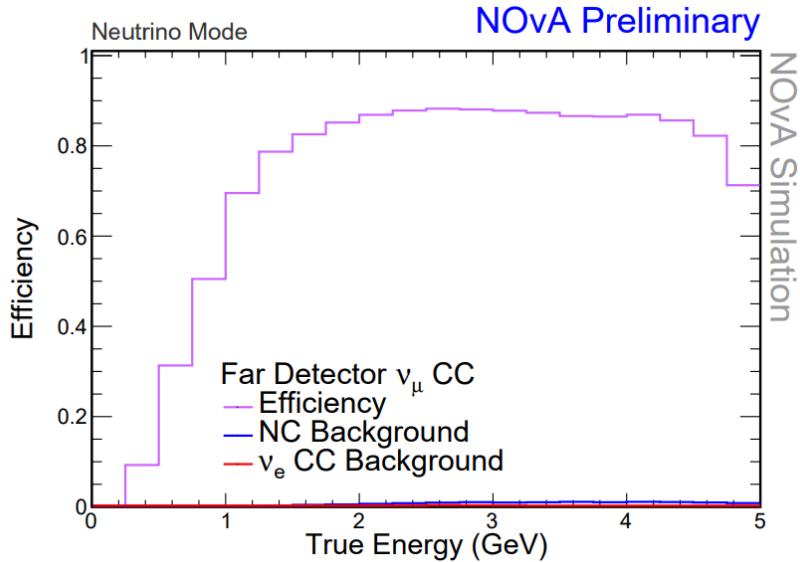


**Figure 6.12:** NOvA’s sensitivity to reject maximal mixing (left) and measure the octant of  $\theta_{23}$  (right) for the normal (top) and inverted (bottom) mass orderings, as a function of the true value of  $\theta_{23}$ . Plots shown for a future  $72 \times 10^{20}$  POT scenario, split evenly between FHC and RHC data. The dotted lines indicate the nominal fake data fits while the solid lines include an additional 3% constraint on  $\Delta m_{32}^2$  around NOvA’s 2020 best-fit value.

## 6.2 Reclaiming $\nu_\mu$ Events

The  $\nu_\mu$  disappearance analysis gives us sensitivity to the values of  $\Delta m_{32}^2$  and  $\sin^2 \theta_{23}$  through measurement of FD  $\nu_\mu$  CC events in the oscillation dip region between 1 and 2 GeV. However, below 2 GeV our efficiency for selecting  $\nu_\mu$  CC events begins to fall off, becoming very low below 1 GeV. This is shown in Fig. 6.13. We were interested in whether this could be improved by reclaiming some of the events that failed selection cuts via a new, secondary selection. This additional sample of reclaimed  $\nu_\mu$  events could then be included in the analysis to hopefully boost the overall selection efficiency, and sensitivity to oscillation parameters.

We began by checking the number and types of events failing cuts. Specifically we

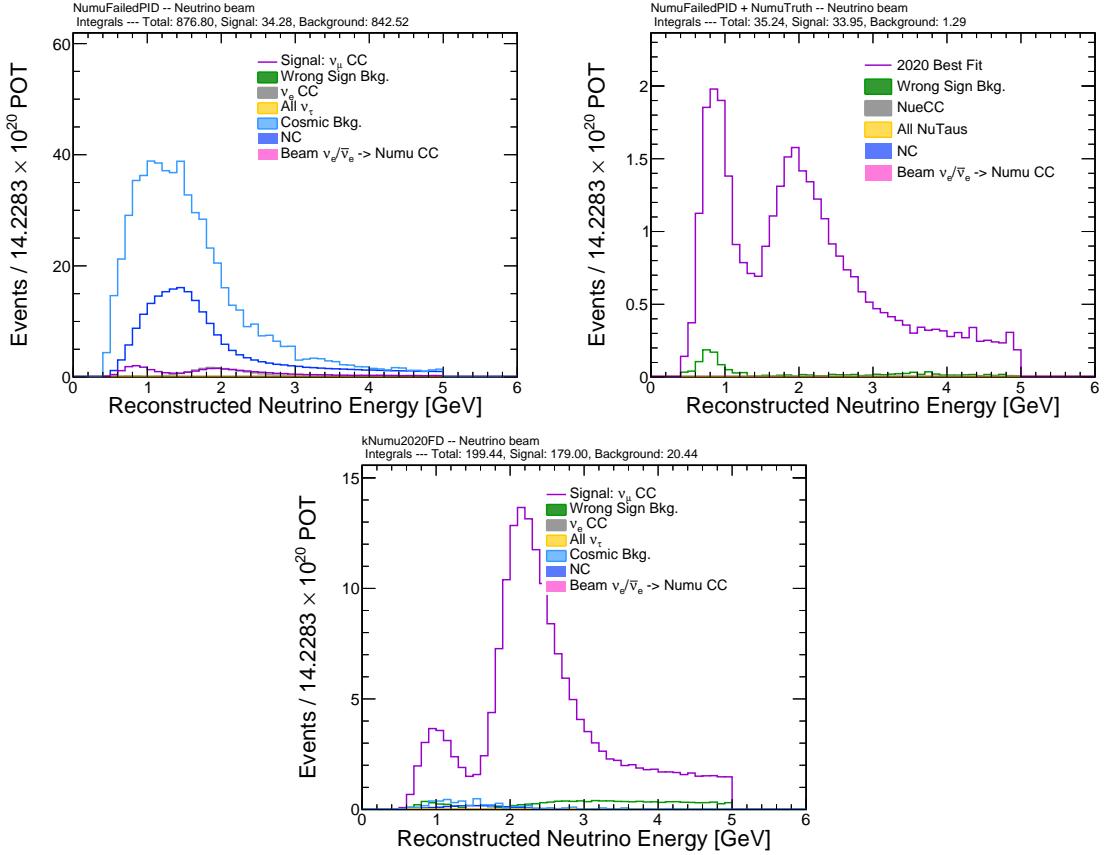


**Figure 6.13:** Selection efficiency as a function of true energy for  $\nu_\mu$  CC signal events, as well as mis-identified NC and  $\nu_e$  CC background, for NOvA’s 2020 analysis. From [152].

looked at the events failing the 2020  $\nu_\mu$  PID selection cut, for the FD FHC simulated dataset. Fig. 6.14 shows the full failed sample, along with specifically the failed signal events we would want to reclaim. These failed signal events would represent a  $\sim 17\%$  increase in the signal of the full selected sample. However, there are large amounts of background in the full failed sample, which are also concentrated at these lower energies. The largest backgrounds are from cosmic and NC events.

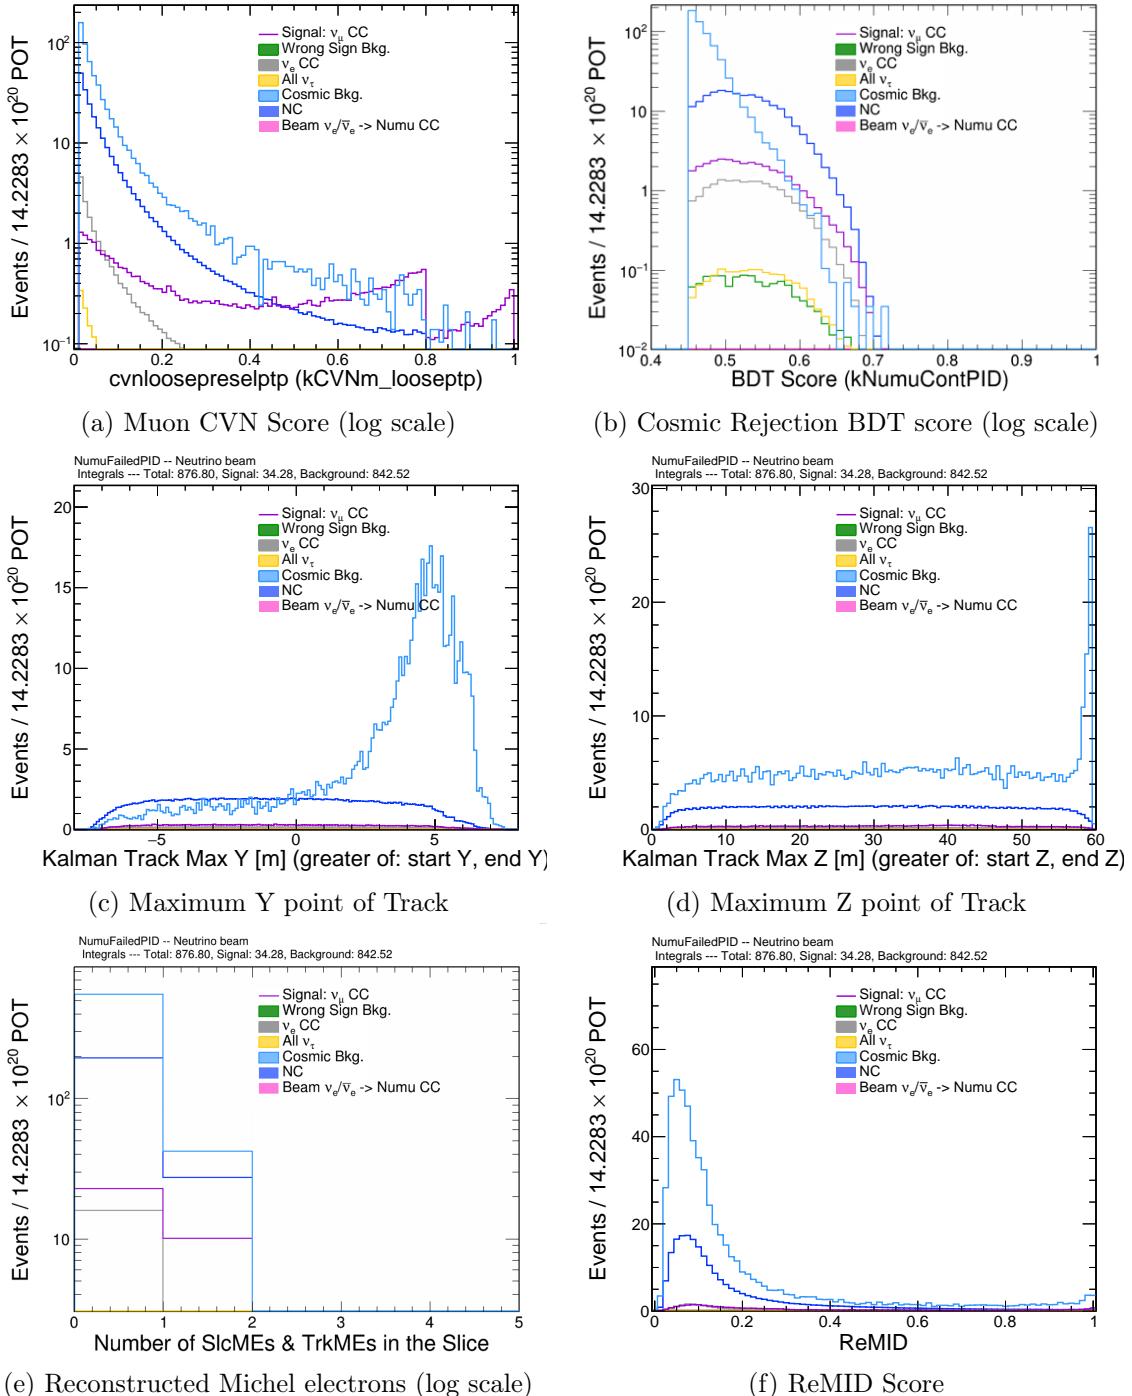
The next step was then to identify additional reconstructed variables we could cut on to reduce these large backgrounds. Fig. 6.15 shows six of the variables we considered that had some potential for separating signal and background. The first variables which stood out were the max Y and Z positions of the muon track, because of the obvious skew in the cosmic background towards higher values. To identify the optimal location to place a cut we calculated the figure of merit

$$FOM = \frac{signal}{\sqrt{background}} . \quad (6.8)$$

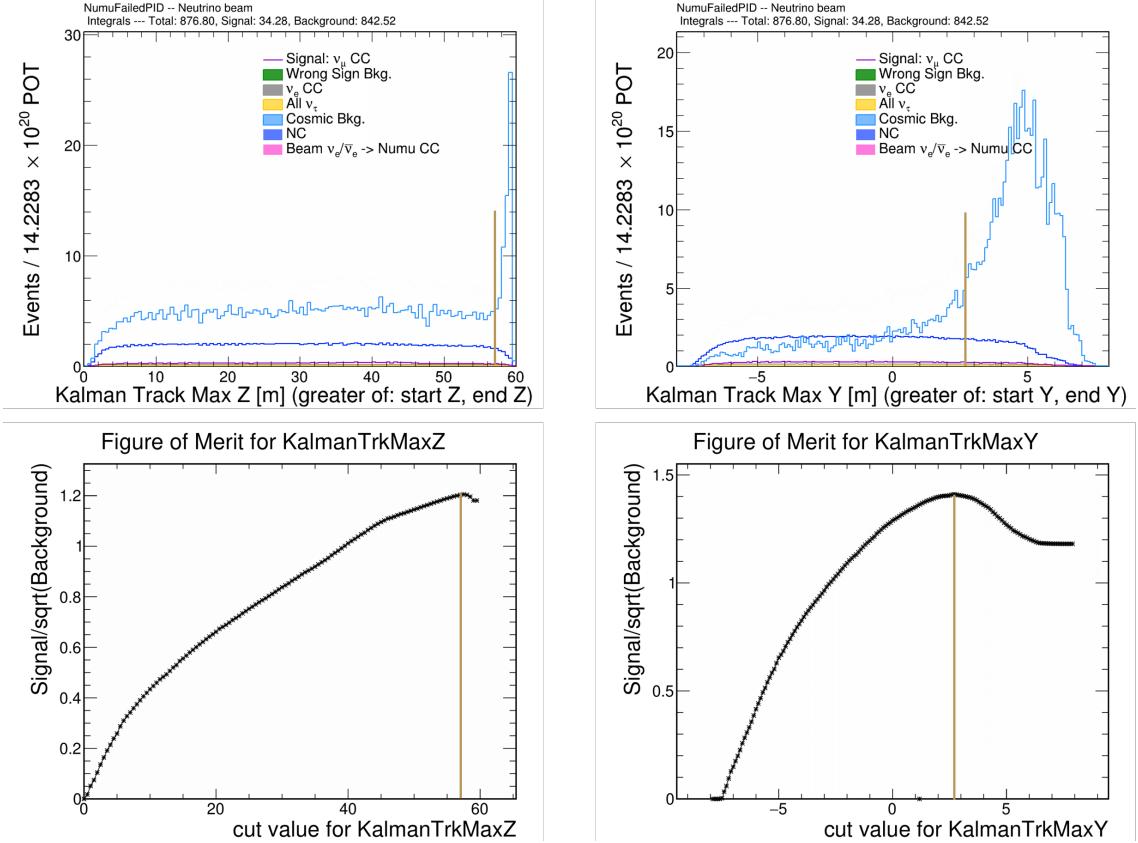


**Figure 6.14: Top Left:** All events failing the 2020  $\nu_\mu$  PID cut, for the FHC sample. **Top right:** All signal events failing the 2020 numu PID cut, for the FHC sample. **Bottom:** All events passing the full 2020 FHC  $\nu_\mu$  selection.

The FOM is plotted in Fig. 6.16 for each position variable, with the optimal cut locations indicated. The effect of applying both optimal cuts to the failed  $\nu_\mu$  sample is shown in Fig. 6.17. The total backgrounds are reduced from 842.52 to 312.90 (62.8% reduction), while the signal goes from 34.28 to 26.33 events (23.2% reduction). The majority of the reduction is in the cosmic background, which is now below the level of the NC background. Despite this the overall background remains high relative to the signal.

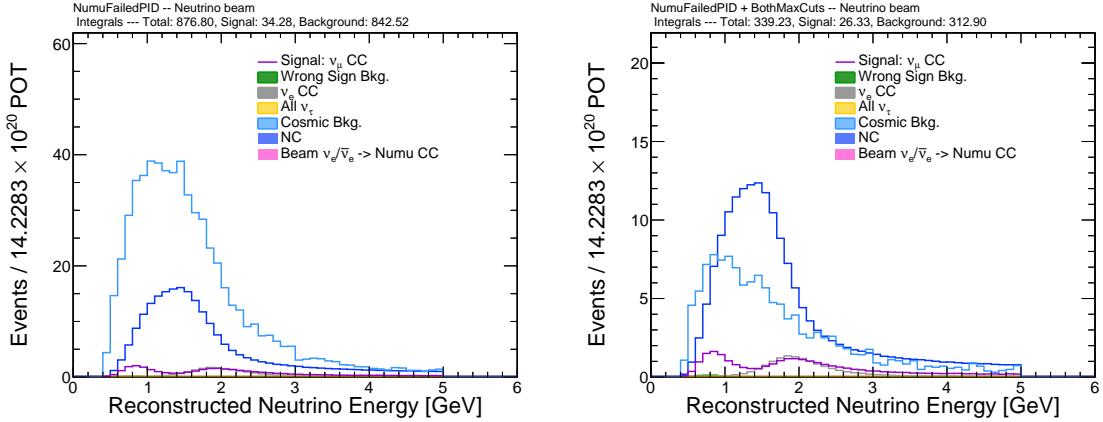


**Figure 6.15:** Plots of simulated events from the 2020 analysis FHC sample failing the  $\nu_\mu$  2020 PID cut.



**Figure 6.16:** The figure of merit for a cut on the max Y and Z positions of the track, for events failing the 2020  $\nu_\mu$  PID selection. The cut removes all events above the current value, and the FOM is calculated with the remaining events. The gold lines indicate the maximum FOM, and corresponding cut locations on each plot.

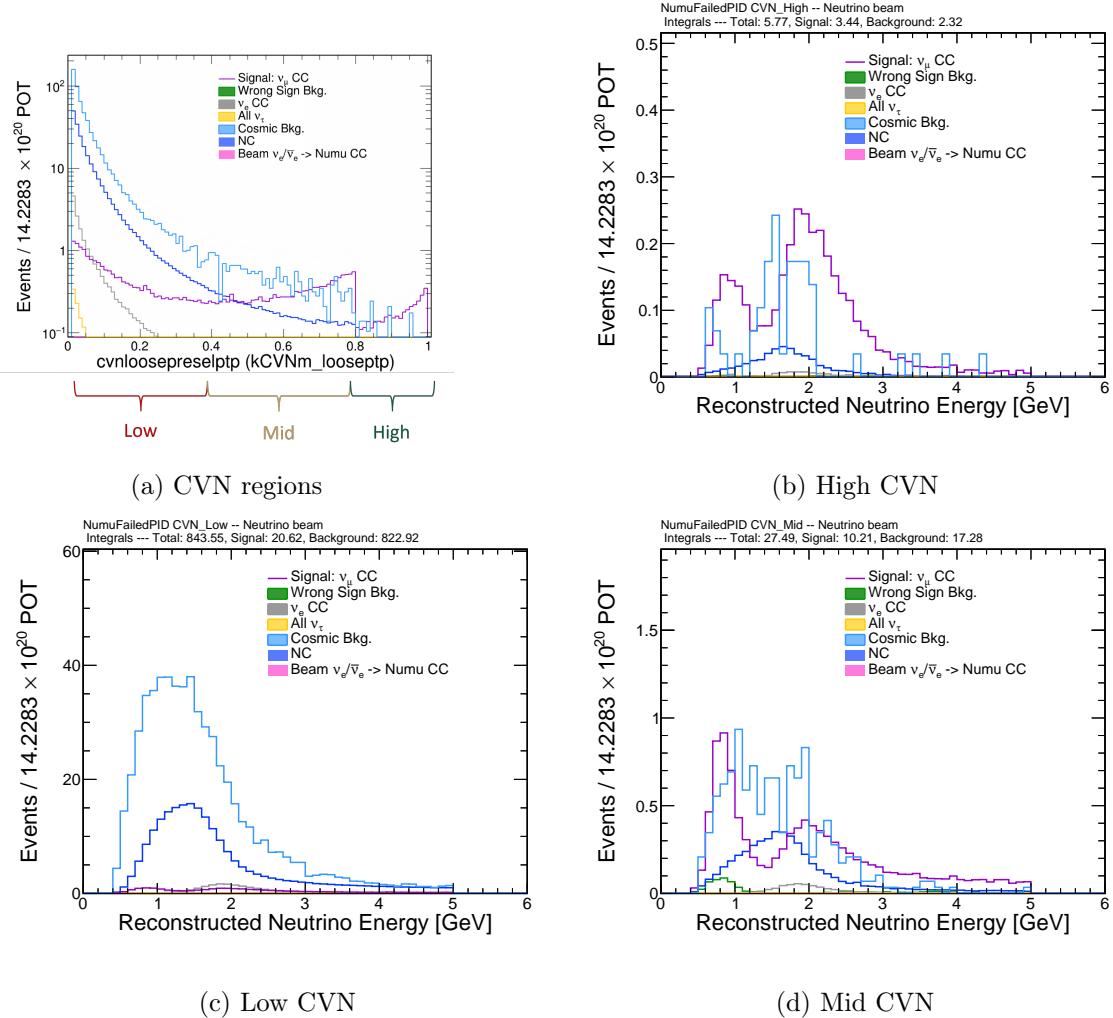
Another variable that showed differences in the signal/background distributions was the CVN score. We considered splitting the failed  $\nu_\mu$  sample by CVN score, similar to what is done with the  $\nu_e$  core sample. We identified 3 different regions in the CVN plot, outlined in Fig. 6.19a. The low region is where  $\text{CVNm} < 0.4$  and corresponds to where the background becomes significantly higher than the signal. The middle region is defined by  $0.4 \leq \text{CVNm} < 0.8$  and corresponds to the region where the signal is approximately equal to the background. Finally, the high CVN region where  $\text{CVNm} \geq 0.8$  is where there is more signal than background. These are events that passed the CVNm cut but failed one of the other PID cuts (cosmic rejection or ReMID), which explains the sharp cutoff in the



**Figure 6.17:** Reconstructed energy distributions for the failed  $\nu_\mu$  event sample before and after cutting on maximum track Y and Z position, using the optimal cut values found in Fig. 6.16.

events at 0.8. The energy distributions for each of these three CVN regions are shown in Fig. 6.18. A large majority of the backgrounds are confined to the low CVN bin, while the mid and high CVN bins have lower backgrounds while retaining just under 40% of the signal events. Additionally, the number of Michel electrons present in the events was higher on average for the signal compared to background, more so for the mid and high CVN bins (shown in Fig. 6.19). Overall, 33% of signal events contained at least one Michel electron, compared to 15% of NC and 7% of cosmic events.

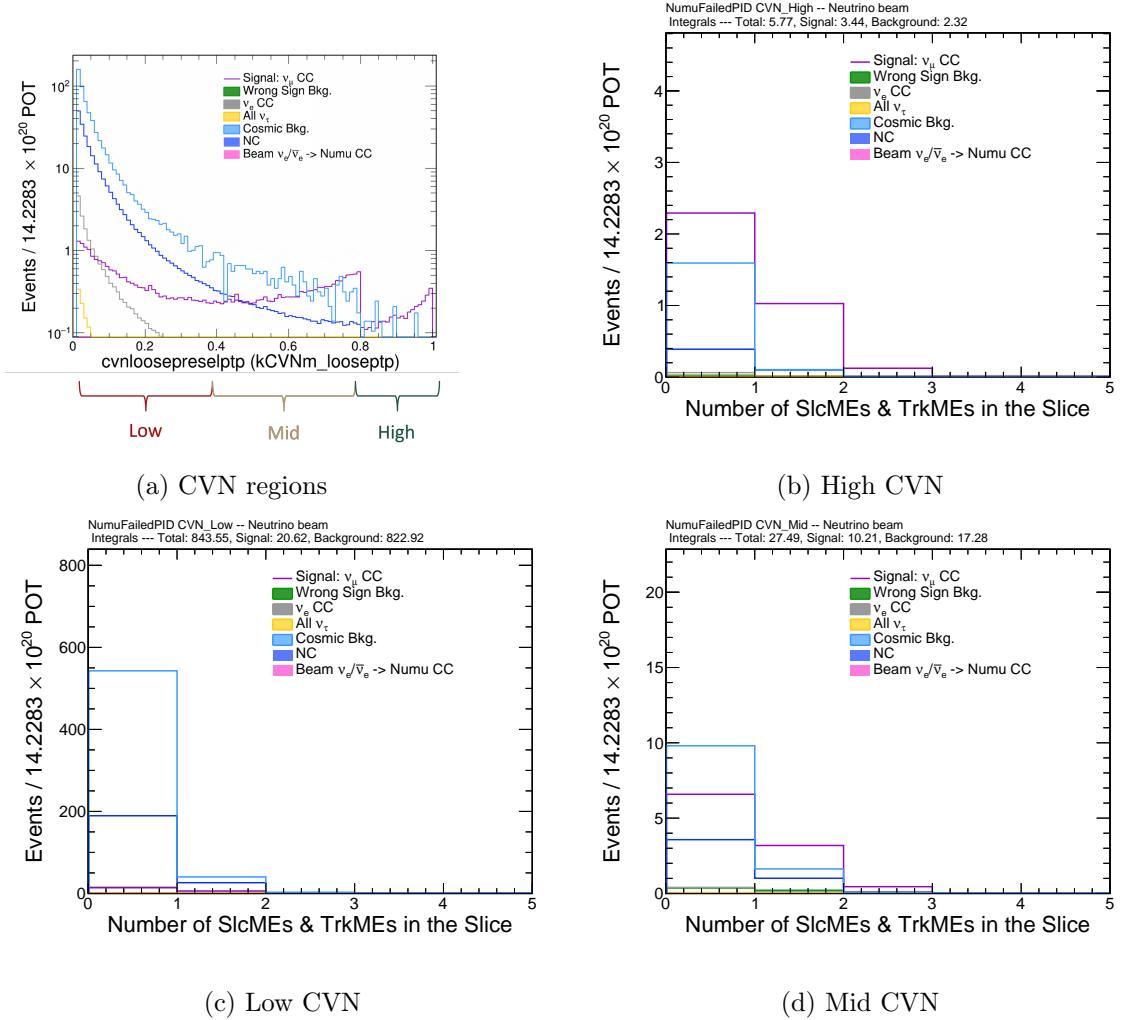
Despite identifying some preliminary variables that could reclaim signal events with reduced backgrounds, a majority of the signal events were still associated with large backgrounds, primarily from cosmic events. Before investing a lot of time in reducing the backgrounds further we wanted to gauge the potential impact of these events on our measurements to see if the effort was warranted. We performed fake data fits to produce 2D confidence regions in  $\Delta m_{32}^2$  vs  $\sin^2 \theta_{23}$ , as well as 1D confidence intervals for each individually. The fake data was generated at the 2020 analysis best-fit point, and only the FHC samples were used in the fitting process, without applying the effects of systematics (stats-only). Three different “best-case” scenarios were considered for the reclaimed  $\nu_\mu$  sample. In all three, the cosmic backgrounds were removed. Then, the first case adds the



**Figure 6.18:** Reconstructed energy distributions for  $\nu_\mu$  events failing the PID cut, split into the 3 bins of CVNm score defined in the top left.

full failed sample (minus cosmics) as an additional sample into the fit. The second case takes the full failed sample (minus cosmics), and splits it into the 3 bins of CVNm score before adding it into the fit. Finally, the last sample takes only the true signal events from the failed sample, and adds them into the fit, representing the theoretical maximum additional sensitivity <sup>2</sup>. The results of the 2D fits are shown in Fig. 6.21, and the 1D fits are shown in Fig. 6.20. The maximum improvement, when adding in only the addi-

<sup>2</sup>We explored splitting the true signal events by CVN score as well, but this didn't offer any improvement over just including them in a single sample.

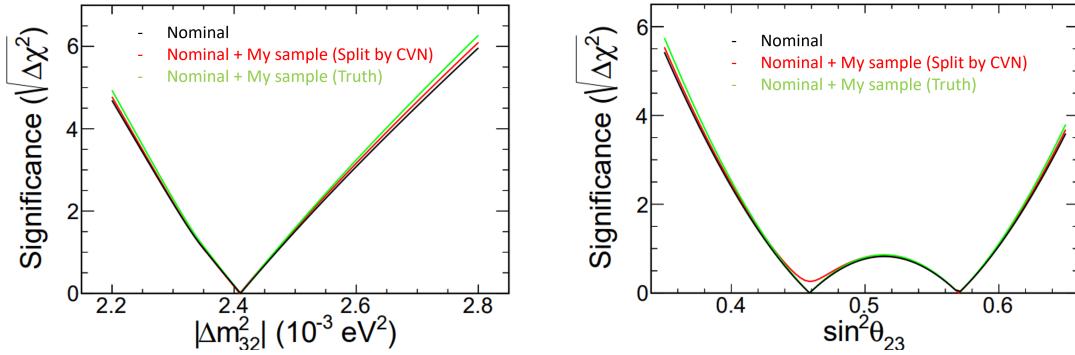


**Figure 6.19:** Number of reconstructed Michel electrons present in the event, for  $\nu_\mu$  events failing the PID cut, split into the 3 bins of CVNm score defined in the top left.

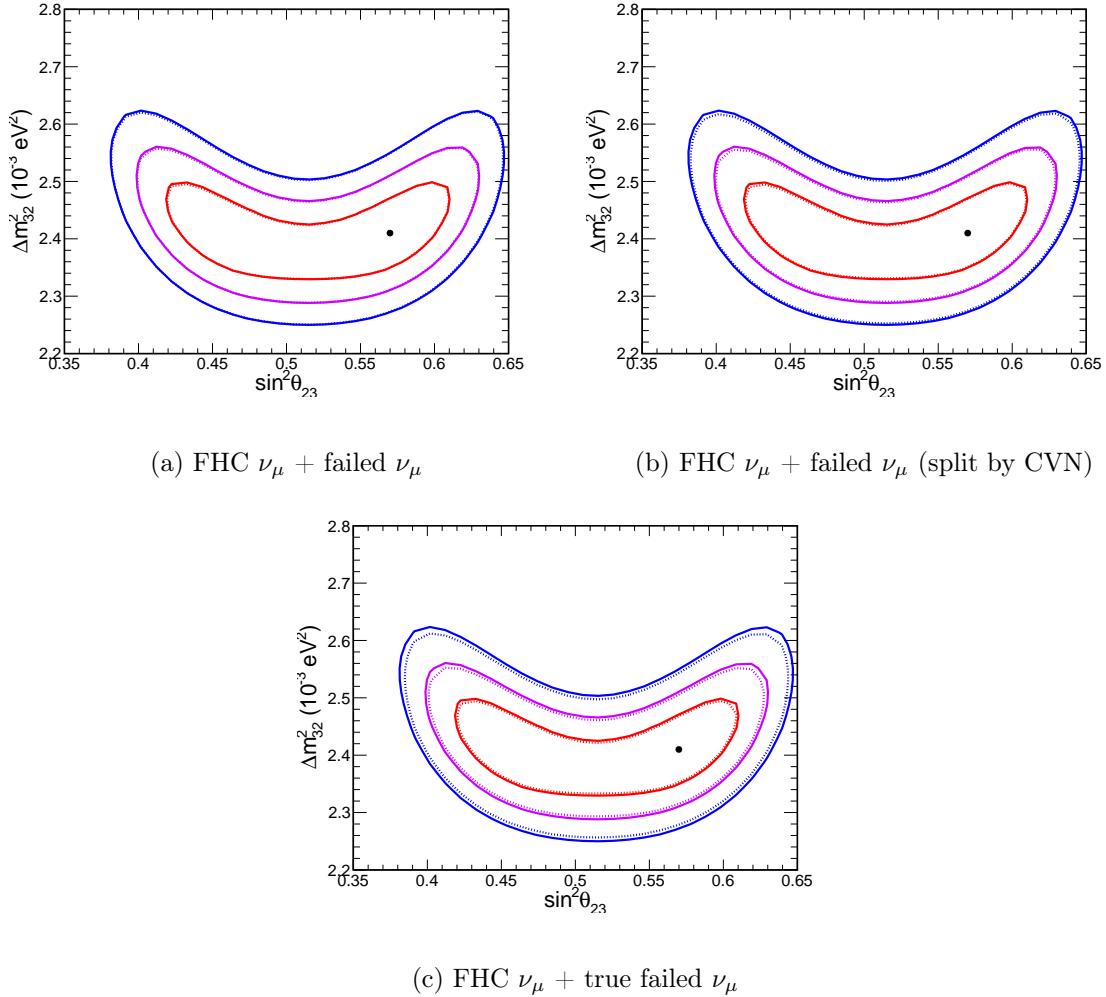
tional true signal events, corresponds to a small but noticeable reduction in the size of the contours. However, this improvement is significantly reduced when instead the full failed sample is included in CVN bins, and barely visible when including it as a single sample. The differences are a bit more obvious in the 1D profiles, where you can see the curves get pulled inwards slightly with the addition of more signal events.

The improvements seen here are very small, and this is a significantly pared-down version of the full fit. Any improvement would likely be washed out by the inclusion of

RHC data. For this reason we elected to not pursue the study further at the time. However, as we near the end of NOvA's datataking, small gains from analysis optimizations will become the main way to improve sensitivity. It may be worth reexamining the study at a future date with the full NOvA dataset, including RHC. We will have taken significantly more data than what was available at the time of the original study, which also means more events available to be reclaimed. Additionally many of the variables will have been re-optimized, so it may be possible to reclaim a purer sample of  $\nu_\mu$  events.



**Figure 6.20:** 1D sensitivity to  $\Delta m_{32}^2$  and  $\sin^2 \theta_{23}$  for a nominal fake data fit, versus 2 fits that include an additional sample of reclaimed  $\nu_\mu$  events. All fits are FHC-only and stats-only, with fake data generated at the NOvA 2020 best-fit point.



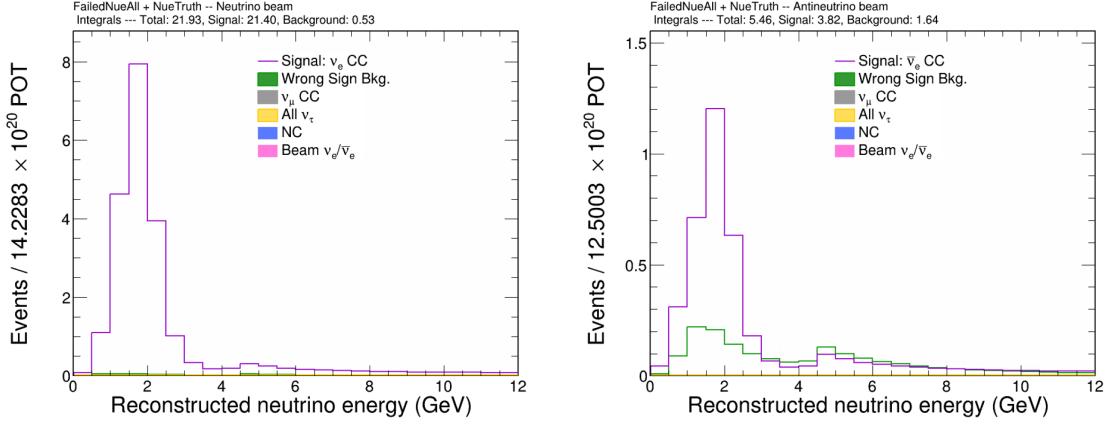
**Figure 6.21:** Fake data fits with/without an additional sample of reclaimed  $\nu_\mu$  events. All fits are FHC-only, and stats-only, with fake data generated at the NOvA 2020 best-fit point. Red, pink, blue lines correspond to  $1\sigma$ ,  $2\sigma$ , and  $3\sigma$  contours. Solid lines indicate the nominal fit without the additional sample, and dotted lines indicate the fits with the additional sample. 3 cases are considered: (a) including the reclaimed events as a single additional sample, (b) splitting the reclaimed events into 3 samples based on CVN score, and (c) including only the true reclaimed signal events as an additional sample. The best-fit point is indicated as the black dot. The additional samples offer a very small but visible improvement in (c), whereas (a) and (b) remain nearly identical.

### 6.3 Reclaiming $\nu_e$ Events

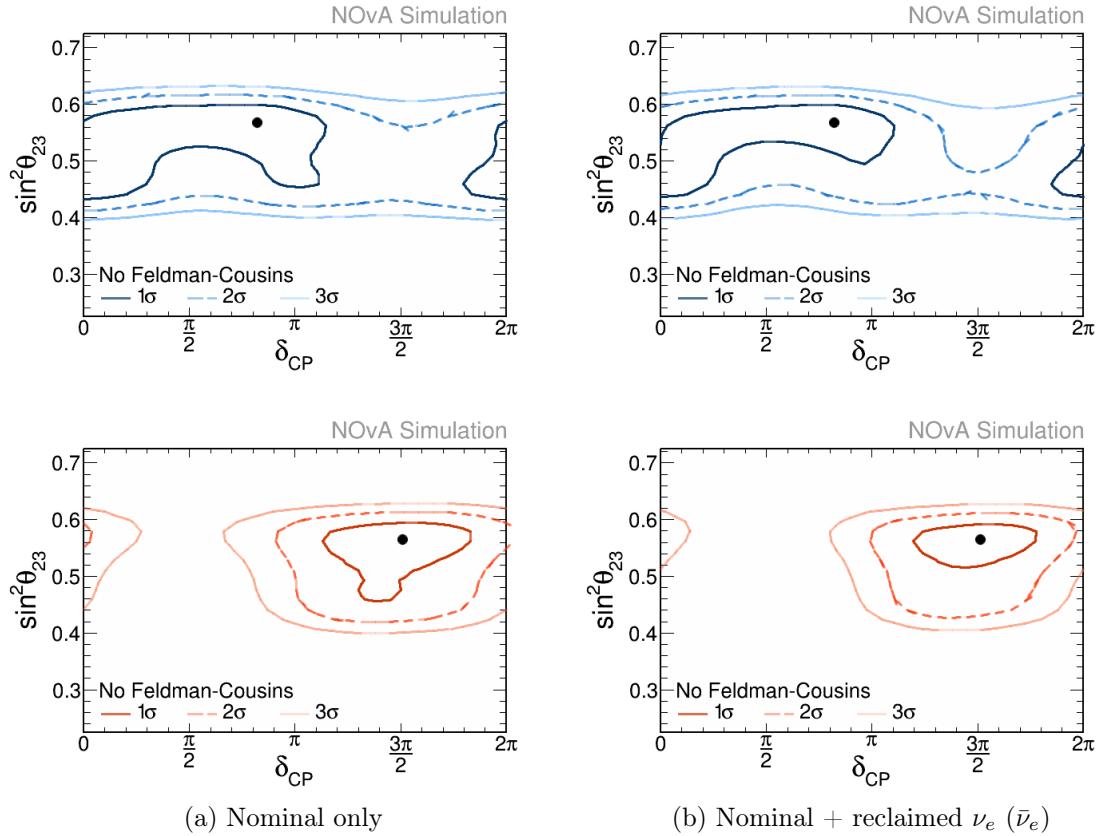
The  $\nu_e$  appearance analysis gives sensitivity to the mass ordering,  $\delta_{\text{CP}}$ , and the octant of  $\theta_{23}$  through measuring differences in the rates of  $\nu_e$  and  $\bar{\nu}_e$  events. As such, knowing the exact counts of each event type is important and effort has been made to incorporate additional samples of  $\nu_e$  events, initially through the peripheral sample, and now for this analysis, the low-energy sample. In the lead-up to this analysis we conducted studies to understand how impactful reclaiming additional  $\nu_e$  events would be, and also where the events were failing cuts, so we could direct efforts to reclaim them. This then served as additional motivation for the low-energy sample, which is discussed at the end of this section.

Starting with NOvA’s 2020 analysis selections and dataset, we looked at the number of  $\nu_e$  events that failed both the core and peripheral sample selections. These  $\nu_e$  ( $\bar{\nu}_e$ ) events are shown in Fig. 6.22 for both the FHC and RHC samples. There are enough events available in these samples to warrant further investigation. If reclaimed, the 21.93 FHC events would represent a 26.7% increase over the 82 events that had been observed for the 2020 analysis  $\nu_e$  data, while the 5.46 RHC events represent a 16.5% increase over the 33 events that had been observed in the  $\bar{\nu}_e$  data [113].

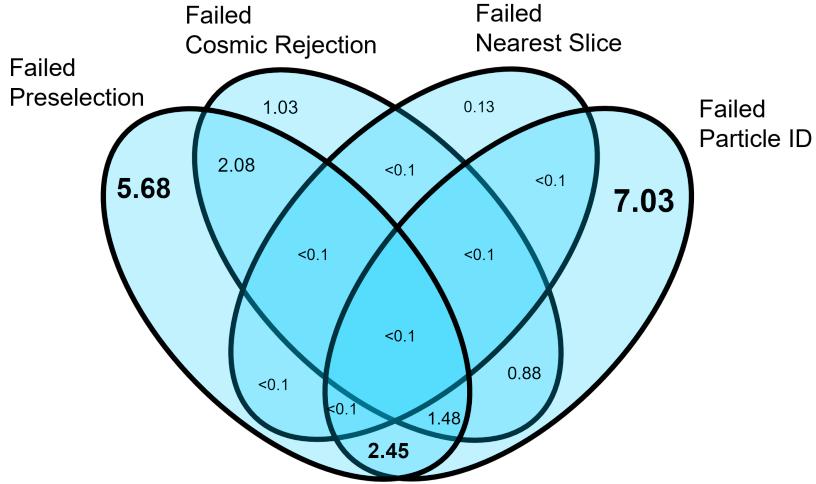
To see what the effect would be of reclaiming all of these events in to our analysis, we performed fake data fits producing confidence regions for  $\sin^2 \theta_{23}$  vs.  $\delta_{\text{CP}}$ , with and without the extra events. The fake data was generated at the NOvA 2020 best-fit point. We include both  $\nu_\mu$  and  $\nu_e$  predictions in the fit, without the effects of systematics (stats-only), and optionally including the additional reclaimed  $\nu_e$  samples (Fig. 6.22). The result is shown in Fig. 6.23. Going from the nominal fit, to nominal with the reclaimed  $\nu_e$  events, we see a large shrinking of the contours indicating that the extra events are enhancing our sensitivity. While this is the “best-case” scenario where we reclaimed all events, it shows that they do have power to constrain the oscillation parameters, and are worth pursuing.



**Figure 6.22:** True FD  $\nu_e$  ( $\bar{\nu}_e$ ) events failing both the 2020  $\nu_e$  core and peripheral selections, for the 2020 analysis FHC (left) and RHC (right) samples.



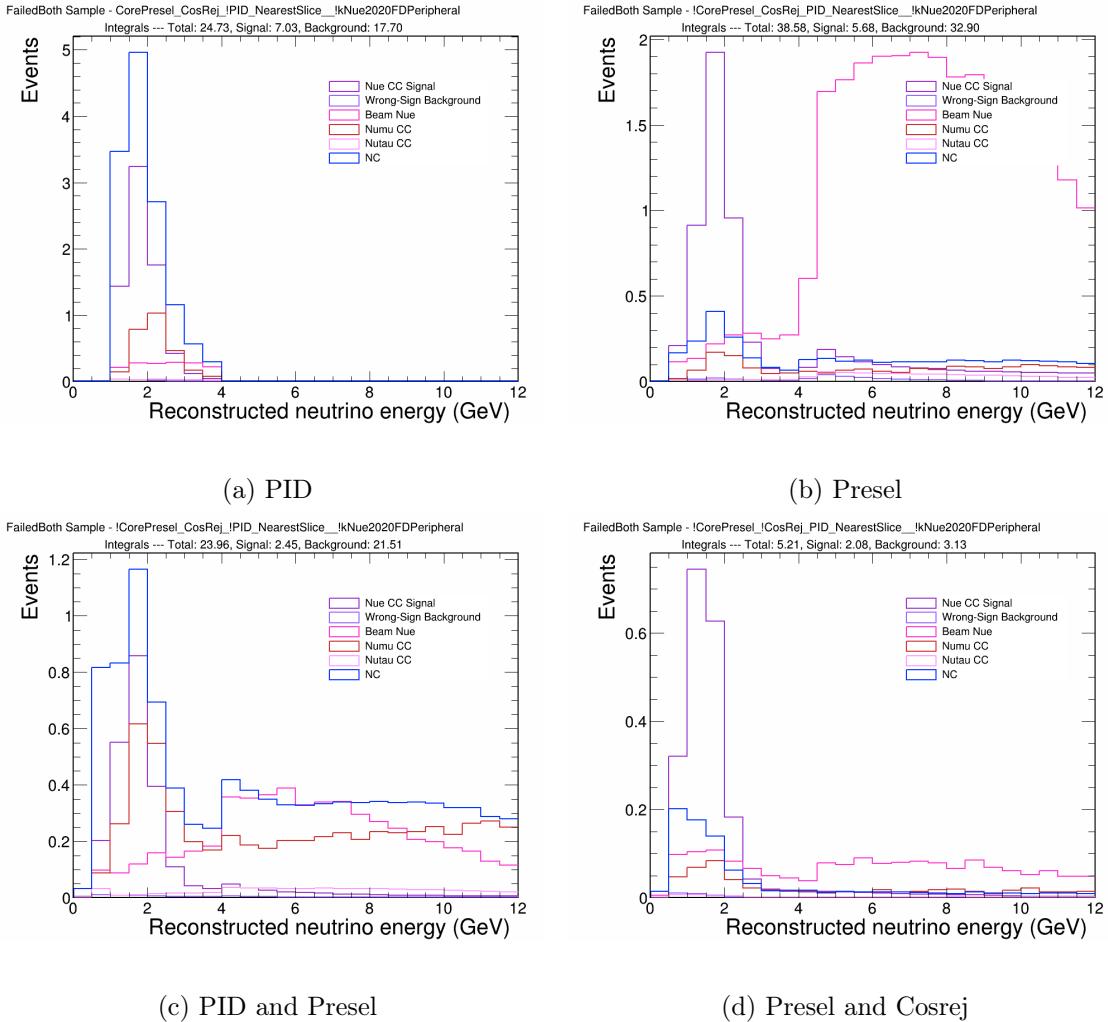
**Figure 6.23:** 2D confidence-level contours for  $\sin^2 \theta_{23}$  vs.  $\delta_{CP}$  in the normal (top) and inverted (bottom) mass orderings, with (left) and without (right) the additional FHC and RHC reclaimed  $\nu_e$  samples. From fits to fake data generated at the NOvA 2020 best-fit point.



**Figure 6.24:** A Venn diagram showing the number of true  $\nu_e$  signal events that failed a particular combination of the 2020  $\nu_e$  selection cuts. Each ring represents one of the 4 cuts making up the 2020  $\nu_e$  selection, with the events in that ring having failed that specific cut.

We then looked into which cuts the events were failing to inform studies attempting to reclaim them. We focused on the FHC sample, since that contained the majority of  $\nu_e$  events we could reclaim, and the 2024 analysis would contain even more FHC data, with the same sample of RHC data as the 2020 analysis. The  $\nu_e$  2020 Core selection has four main cuts, Preselection, CosRej, PID, and NearestSlice. These are similar to the cuts described in Sec. 5.2.2, but organized differently. Preselection covers reconstruction and event quality, and containment cuts. The PID cut here is just a cut on CVNe score, while CosRej contains the cut on cosmic rejection BDT score. NearestSlice is the same cut as in Sec. 5.2.2. We constructed predictions for events failing specific combinations of these cuts, and removed any events that would have been selected into the Peripheral sample. We then counted up the number of events that fall into each category. For example, we found that 7.03 of the 21.93 total events failed only the PID cut, while passing the other cuts. There are 2.45 events failing both PID and Preselection cuts, while passing the other cuts, and so forth. We can organize all of the categories of failed events into a Venn diagram which shows the number of true signal events in each category. This

is shown in Fig. 6.24. Most of the events are failing just the PID cut (7.03 events), or just the Preselection cut (5.68 events). After these two categories, the next-largest are events failing both PID and Preselection (2.45 events) and events failing both Preselection and CosRej (2.08 events). The remaining  $\sim 3$  signal events are spread between the other combinations of PID, Preselection, and CosRej. The NearestSlice cut is removing a very specific class of background event so the categories with it have only a small amount of signal.

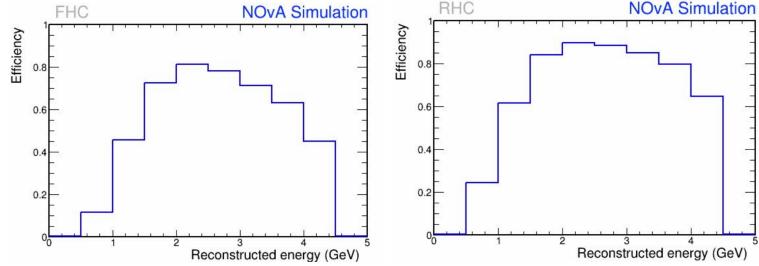


**Figure 6.25:**  $\nu_e$  signal and beam background energy distributions, for the 4 Venn diagram categories containing the most  $\nu_e$  signal. Labels indicate the cut categories that were failed.

Next, we examined the beam backgrounds. Fig. 6.25 shows the energy distributions of both the signal and beam backgrounds for the four largest Venn diagram categories. The largest beam backgrounds in our energy region of interest ( $< 4$  GeV) are present in the events which failed the PID cut, and are comprised of mostly NC events, with some mis-identified  $\nu_\mu$  CC events as well. There are beam  $\nu_e$  events concentrated at higher energies, which we see mostly in the events failing Preselection, since that cut selects neutrino energies between 1 and 4 GeV. The plots do not include cosmic events, but we would expect some cosmic background in the samples that failed Preselection (due to its containment cut) and larger cosmic backgrounds in the samples that failed the CosRej cut. From these distributions we learned that if we are able to reclaim these  $\nu_e$  events we should expect larger NC backgrounds in the sample, and depending on the category of cut we target, there may be additional cosmic backgrounds as well.

Given the large backgrounds, it was not feasible to reclaim events by simply retuning our existing selections, and a more sophisticated approach would be needed. One avenue to reclaim events that we pursued was the low-energy BDT (described in Sec. 4.2.5). This targeted  $\nu_e$  events in the 0.5 - 1.5 GeV range, and could potentially reclaim events from several of the categories we examined. As explained in Sec. 5.2.2 the low-energy selection allows a wider range of CVN score, so it can target  $\nu_e$  events failing the PID cut. Additionally, the energy range allows us to include events falling below the Preselection energy cut.

There were several additional reasons for the focus on lower neutrino energies. First, similar to the  $\nu_\mu$  selection, the  $\nu_e$  selection efficiency also falls off at lower energy (shown in Fig. 6.26). This is an indication that there are events we could reclaim, but also that they may be difficult to identify using our nominal selections, another motivation for training a new BDT explicitly on low-energy events. Additionally, at lower energies we have a higher sensitivity to the asymmetry in the rates of  $\nu_e$  ( $\bar{\nu}_e$ ) events. We define asymmetry as the difference in the probability of  $\nu_\mu \rightarrow \nu_e$  and  $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$  oscillations, divided by the total



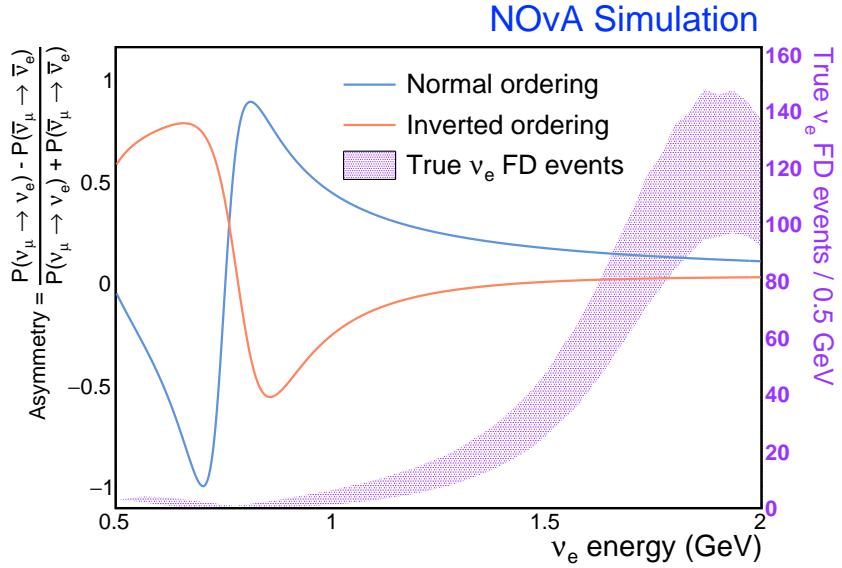
**Figure 6.26:** The efficiency with which we select true  $\nu_e$  CC events in the FD as a function of energy, for FHC (left) and RHC (right) samples. Efficiency is defined as (selected  $\nu_e$ )/(total true  $\nu_e$ ) events in each bin. The NOvA 2020 analysis  $\nu_e$  selection was used, with the Preselection energy cut lowered from 1.0 GeV to 0.5 GeV to include the low-energy region. From [103].

probability:

$$\text{Asymmetry} = \frac{P(\nu_\mu \rightarrow \nu_e) - P(\bar{\nu}_\mu \rightarrow \bar{\nu}_e)}{P(\nu_\mu \rightarrow \nu_e) + P(\bar{\nu}_\mu \rightarrow \bar{\nu}_e)}. \quad (6.9)$$

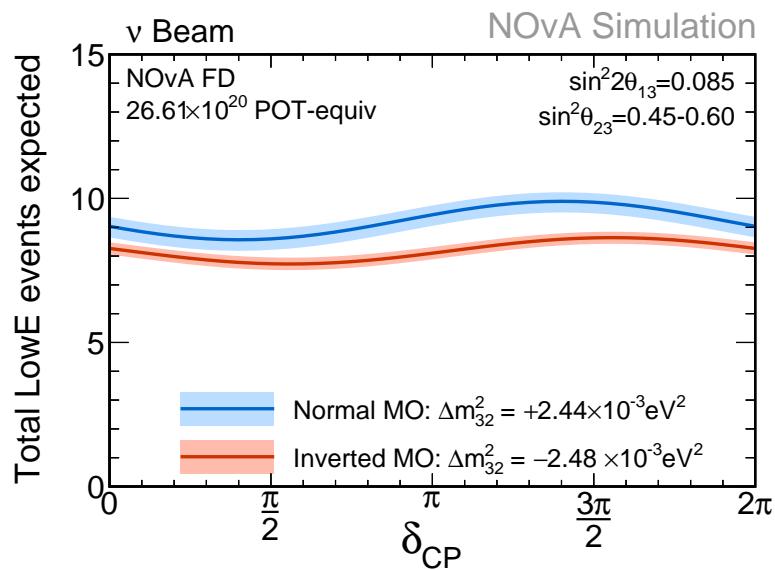
This is plotted as a function of energy in Fig. 6.27, along with the expected range of  $\nu_e$  events. We see that the maximum and minimums occur below 1 GeV, and are inverted for each mass ordering. While we expect significantly fewer events at low energy, they should also be highly sensitive to the choice of mass ordering.

For this analysis we only included the FHC low-energy sample. This is because the corresponding RHC sample was found to be too small, with the total FD prediction size being less than a single event, and the majority of that being background [103]. If NOvA is able to take more RHC data in the future, the sample may become a viable addition to the analysis. Of course, without the RHC sample to compare to, some of the sensitivity to the asymmetry is lost. Despite that, the FHC low-energy sample still maintains some small ability to distinguish between the mass orderings. Fig. 6.28 shows the expected total number of events in the low-energy sample for this analysis as a function of  $\delta_{CP}$ , for each choice of mass ordering. The shaded bands show the variation over the  $1\sigma$  range of  $\sin^2 \theta_{23}$ . On average there is about 1 event separation between the choice of mass orderings. Including the  $1\sigma$  range of  $\sin^2 \theta_{23}$ , the difference between the maximum in the normal ordering and minimum in the inverted ordering is about 3 events.

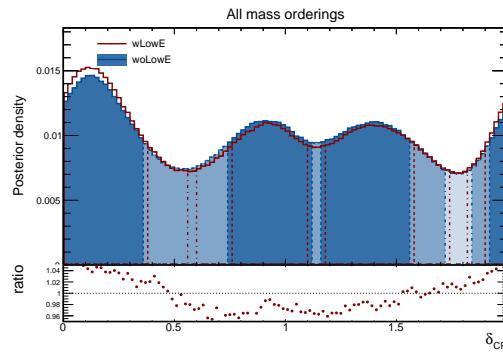


**Figure 6.27:** Predicted asymmetry of FD  $\nu_e$  and  $\bar{\nu}_e$  events at the far detector. Overlaid is a spectrum of predicted  $\nu_e$  events at the FD without any cuts. Asymmetry curves use the best-fit parameters found in Chapter 7. The predicted spectrum spans credible oscillation parameters around these best fit values. From [153].

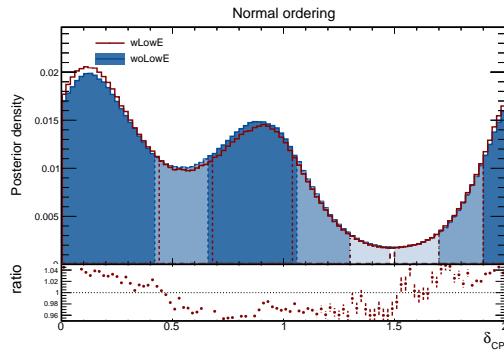
To gauge the impact of the low-energy sample on the full analysis, a fake data fit using the Bayesian framework was conducted both with and without this additional sample. Fits included all systematics, and used the 1D Daya Bay reactor constraint. Fig. 6.29 shows the posterior probability density for  $\delta_{CP}$  from the full fits with and without the low-energy sample, and for 3 different mass ordering cases. The first case allows either mass ordering in the fit, while the second and third are restricted to either the normal or inverted mass orderings respectively. Within the fit that allows both mass orderings (Fig. 6.29a), including the low-energy sample gives a modest increase to the percent of the posterior density in the normal mass ordering from 73.5% to 74.1%.



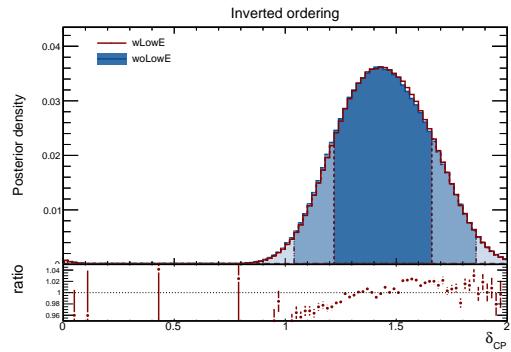
**Figure 6.28:** Total FHC low-energy  $\nu_e$  appearance events predicted as a function of  $\delta_{\text{CP}}$ . The solid lines represent the best-fit value of  $\sin^2 \theta_{23} = 0.54$ , while the band shows the variation over a wider range of  $\sin^2 \theta_{23}$  values listed in the top right.



(a) Both MO



(b) Normal MO only



(c) Inverted MO only

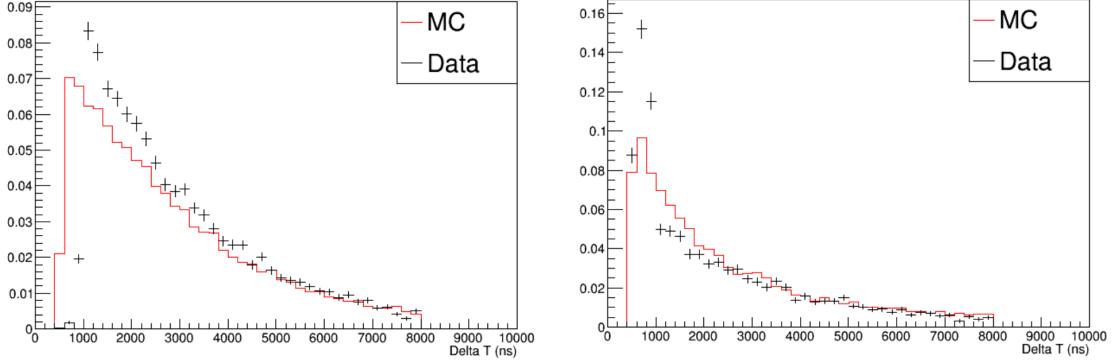
**Figure 6.29:** Comparison of 1D posterior density distributions for  $\delta_{CP}$  from fits with and without the low-energy sample included. For (a) both mass orderings, (b) normal mass ordering only, and (c) inverted mass ordering only cases. The fit is done with all systematics and 1D Daya Bay reactor constraint applied. From [143].

## 6.4 Michel electrons

A large focus of my work has been on improving the simulation and reconstruction of Michel electrons in NOvA. This work was started after production had already begun on the files that would be used for this analysis (referred to as Production 5.1, or Prod5.1 files). Since the simulation and reconstruction code is frozen before starting production, the Michel improvements will be targeted to future analyses. In the interim period, the improvements were tested using a few different sets of files. Initially we produced dedicated sets of files from scratch running the necessary simulation and reconstruction jobs with the improvements added in. Recently, we have been gearing up for a new production run (Production 6) which will include the improvements, and we have been able to begin testing their full implementation. Preliminary samples of files were produced in a “mini-production”, to test and validate the changes before undertaking the full production. Most plots in this section were made from these sets of files.

### 6.4.1 Improving Simulation

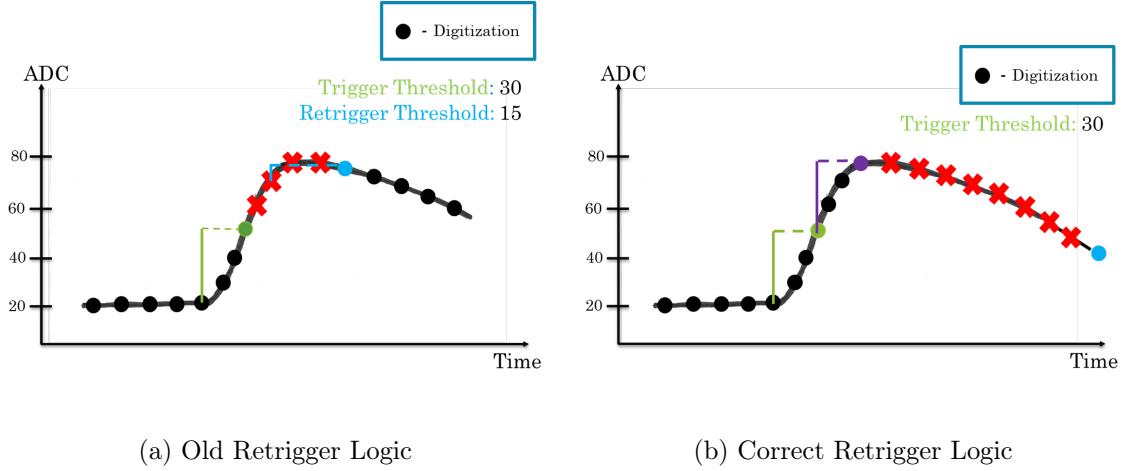
One of the first improvements made was a fix to a longstanding data/MC discrepancy in the  $\Delta T$  of the Michel electrons, the time difference between the Michel and parent slice (nominally a muon). The discrepancy is shown in Fig. 6.30, and is different depending on whether the Michel retriggered one of the cells in the parent slice. It is common for reconstructed Michel electrons to contain hits that retrigger cells from their parent slice, since they are a time-delayed signal that appears close to where the muon decayed. What we saw was an excess in MC at low  $\Delta T$  for Michels that retrigger. For Michels that don’t retrigger we saw the opposite, an excess in data at low  $\Delta T$ . One way this could occur is if the APD deadtime was shorter in simulation than in the actual hardware. Careful examination of the readout simulation and detector firmware code found this to be the case. Fig. 6.31 illustrates the differences in the retrigerring logic between the old simulation, and the real detector firmware. In both cases a dual-correlated sample (DCS) is first used to



**Figure 6.30:** ND Data/MC discrepancy seen in the  $\Delta T$  of reconstructed Michel electrons in older simulation files. Left is for Michels which retrigger a cell from the parent slice, and right is for Michels which do not retrigger. Histograms have been area normalized to illustrate the shape differences. From [154].

trigger a readout. The current ADC value is compared to the one from 3 samples ago, and if their difference is larger than the trigger threshold, we trigger and read out a hit. The old simulation used a variable deadtime, which lasted until the current DCS value fell below a defined “retrigger threshold”. In the real detectors, the peak is first found after triggering (the purple point), and then a fixed deadtime lasting 9 digitizations is implemented. The two detectors have different timing resolutions, so this results in a  $1.125 \mu\text{s}$  deadtime at the ND, and a  $4.5 \mu\text{s}$  deadtime at the FD. Since the deadtime of the old simulation could vary, it often ended earlier than the fixed deadtime of the data (the blue point in the example). Simulated Michel electrons, which appeared soon after the muon decay, could then retrigger cells before the data, hence the discrepancy.

We rewrote the readout simulation code to match the logic from the firmware in each detector, and then produced new simulated files with the change. We then ran the reconstruction code up to MEFinder, and plotted the  $\Delta T$  in MC vs. the data. The result for the ND is shown in Fig. 6.32. The simulation now matches the data and we see the  $\Delta T$  for the retriggered Michels peak and then fall off at the deadtime of 1125 ns as expected.

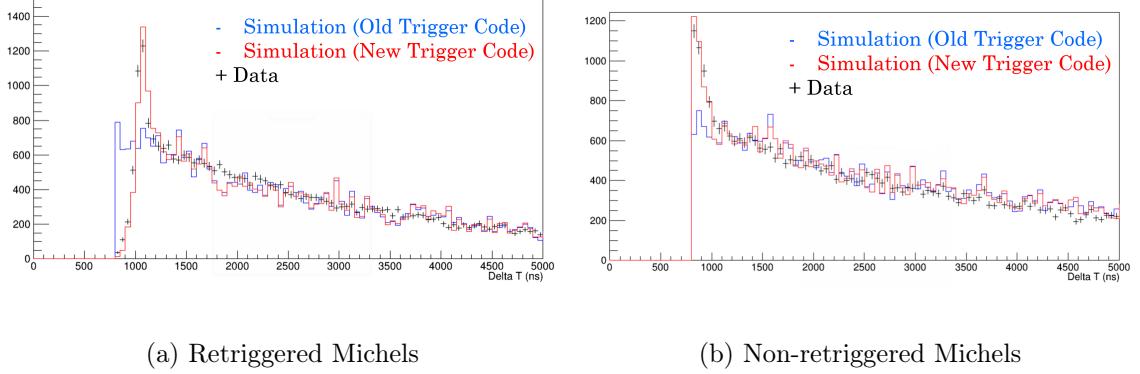


**Figure 6.31:** Cartoon diagrams of the retriggering logic used in the old simulation, and the correct version from the real detector firmware. **Left:** In the old simulation, a trigger was initiated once the DCS value goes above the trigger threshold (green point), and you cannot trigger again until it goes back below the retrigger threshold (blue point). **Right:** In the actual detector, a trigger is also initiated when the DCS value goes above threshold (green point), the peak is then found (purple point), and a fixed deadtime of 9 digitizations is always used (live again on the 9th).

#### 6.4.2 Improving Reconstruction

With the timing simulation fixed, I moved on to examine the Michel reconstruction code, MEFinder (detailed in Sec. 4.1.6). There were several improvements made. With the timing fixed, we removed the cut on retrigged candidate Michel hits from the hit-finding step in MEFinder. We also identified and fixed a bug in the hit-finding that was missing candidate Michel hits at the ends of the 10,000 ns window after a physics slice. The code was attempting to use a binary search function to identify the first and last valid hits in the window, but it was not implemented correctly. Switching to a simple linear search function resulted in no missed hits, and by removing some unnecessary function calls we were able to mitigate any loss in search algorithm efficiency, with the new jobs finishing  $\sim 20\%$  faster [155].

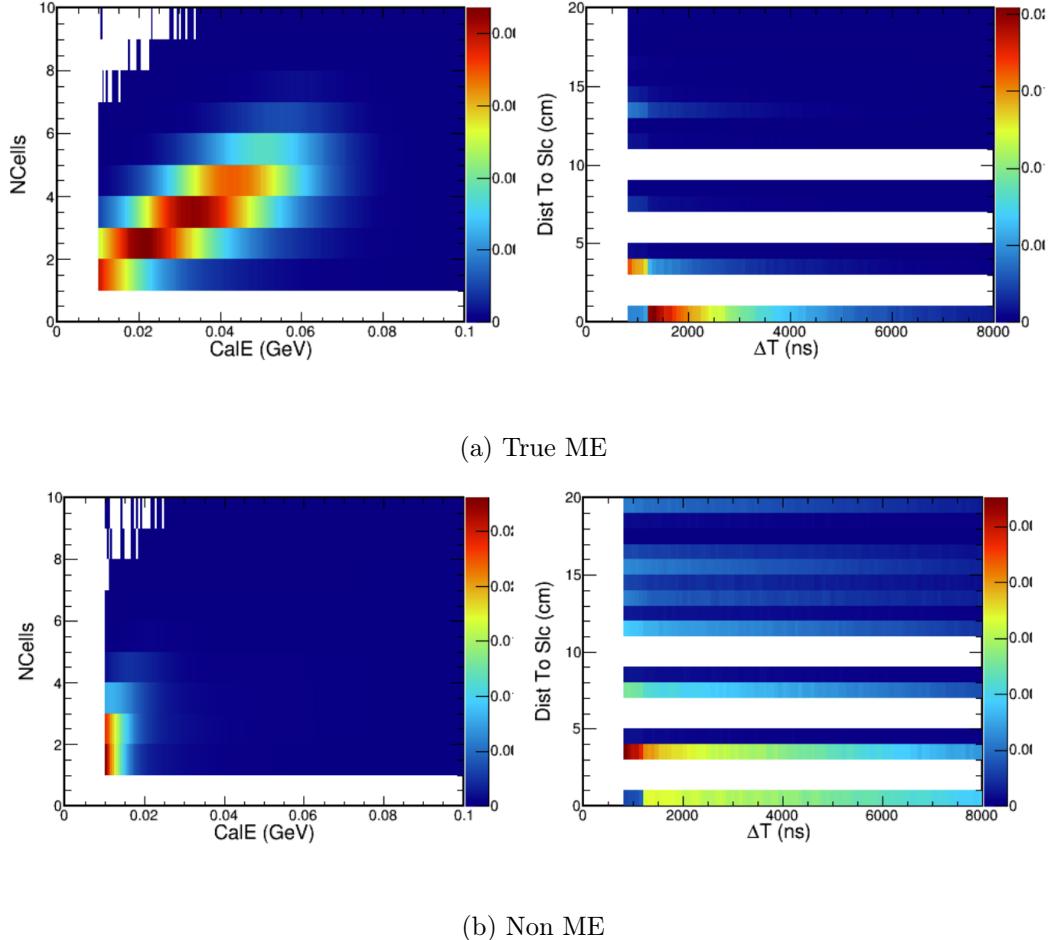
Since the timing distribution had significantly changed, the next step was to remake the MID template histograms (described in Sec. 4.1.6). This was overdue, as they had not been remade in several years despite other changes to the reconstruction code that could



**Figure 6.32:** Data/MC comparison of Michel electron  $\Delta T$  with old and new retrigerring logic for Michels that retrigger (left) and those that don't (right).

have affected the Michels, such as an update to the slicer. Additionally, we identified a reconstruction failure associated with the old templates. The old templates (Fig. 6.33) have a cutoff in the  $\Delta T$  variable at 800 ns. This was put in place early on because of the poor data/MC agreement. However, the cutoff extends to all values of the Distance To Slice variable, instead of only being applied to the 0 DistToSlc bin which contains all the retriggered Michels. As described in Sec. 4.1.6, the true Michel DistToSlc vs.  $\Delta T$  template is used to match reconstructed Michel electrons to their parent slice. So, the knock-on effect of this cutoff is that MEFinder would not match any Michels to parent slices with a  $\Delta T < 800$  ns, even for Michel electrons that didn't retrigger a cell. Instead, they would either be matched to a different parent slice at a later  $\Delta T$  (potentially incorrectly) or, they would not be reconstructed at all. This effect can be seen as the sharp cutoff in events below 800 ns in the  $\Delta T$  plots in Fig. 6.32.

Remaking the MID templates is an iterative process. Since the templates are used in the reconstruction, the initial round of reconstruction with the new code will still have to use the old templates. Then, we can create new template histograms from the newly-reconstructed Michels. We then re-run MEFinder over the same input files, this time using the new templates, to make sure the effects of any changes are fully captured. This final sample of Michels can then be used to create a final version of the template histograms

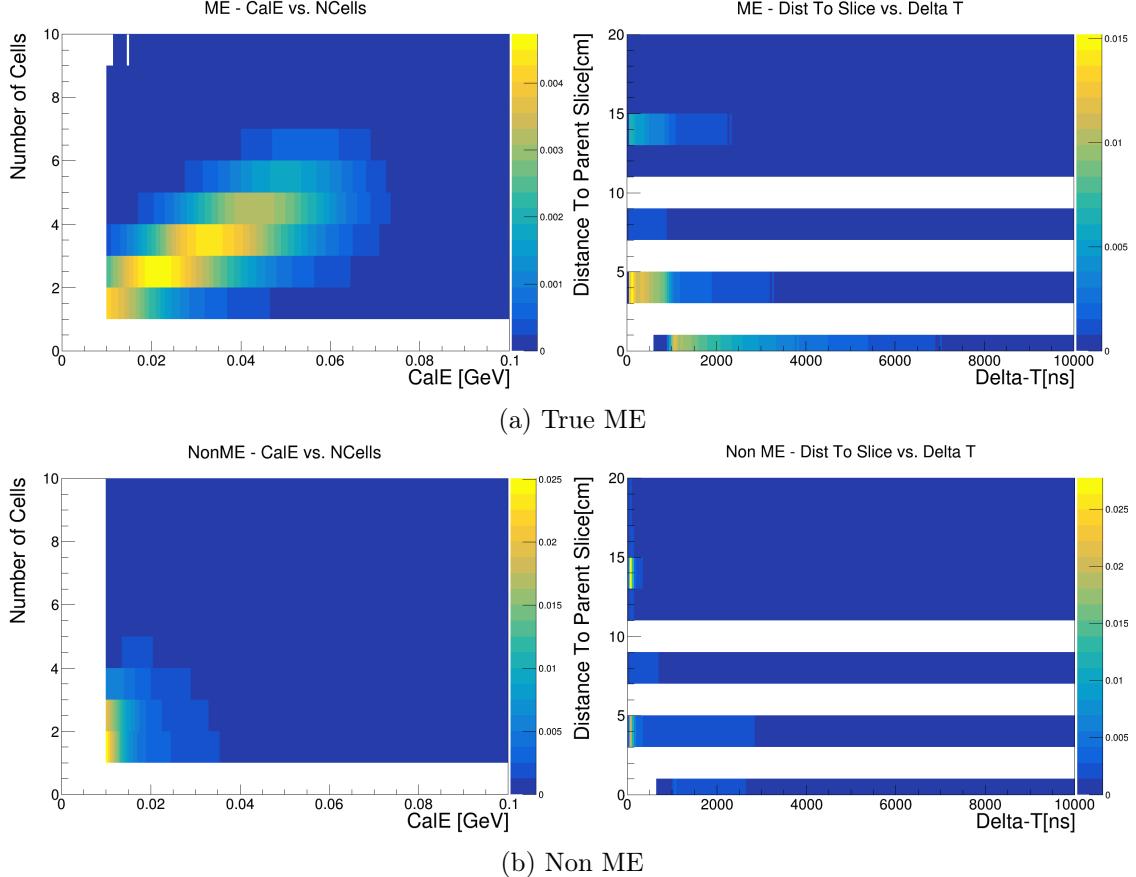


**Figure 6.33:** MID template histograms using the old simulation. The top two plots are for the true-ME case, and the bottom two are for the non-ME case. From [95].

that will be used going forward<sup>3</sup>. There were several rounds of file production in the mini-production leading up to Production 6, which allowed us to do this iterative process over a large enough dataset. The final template histograms incorporating all the changes are shown in Fig. 6.34. The CalE vs. NCell templates are largely the same as before, but there are significant changes to the  $\Delta T$  templates. The cut at 800 ns  $\Delta T$  was removed and we can see there are a significant number of events at low  $\Delta T$ , with the majority of

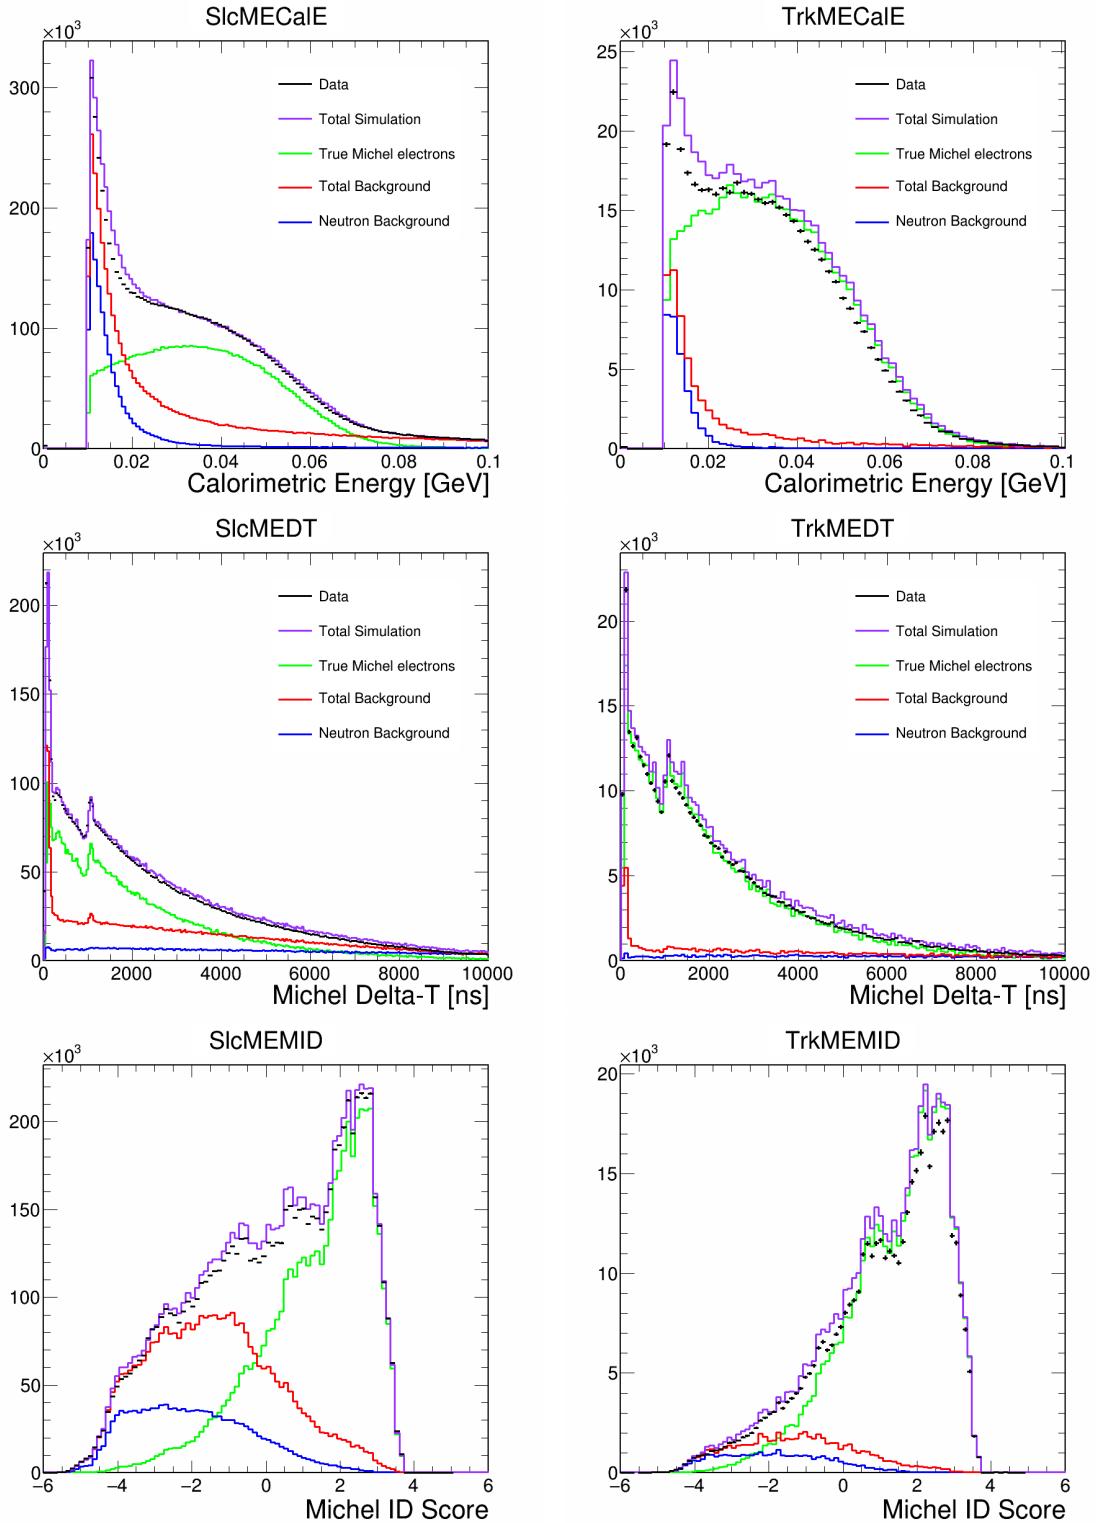
<sup>3</sup>In principle, if there are still differences in the templates or reconstructed Michel distributions after the second iteration, this process could be continued until the distributions no longer change. In practice however the two iterations of reconstruction were sufficient.

the backgrounds occurring at very low  $\Delta T$ .



**Figure 6.34:** ND MID template histograms using the new simulation. The top two plots are for the true-ME case, and the bottom two are for the non-ME case.

A Data/MC comparison of the energy,  $\Delta T$  and MID distributions of Michels using the new MID templates are shown in Fig. 6.35.



**Figure 6.35:** Reconstructed energy,  $\Delta T$ , and MID distributions for SlcME (left) and TrkME (right) in the near detector. These Michels were reconstructed using the ND MID templates in Fig. 6.34.

The shape of the distributions agree well with data. The Gaussian shape of the energy distribution we measure for true Michels can be seen at higher energy<sup>4</sup>, with a spike at lower energies from the backgrounds. The  $\Delta T$  distribution extends all the way to 0 now, and we see a spike around 1125 ns from Michel electrons that retrigger. The MID distributions show most of the signal events being skewed to  $\text{MID} > 0$ , so we can still expect to effectively cut on that variable to remove backgrounds. The MC events have been scaled to match the POT of the data events, but there is still a slight normalization offset between data and MC in several of the distributions, particularly for the TrkME. This is expected to occur because we do not yet have cross-section weights for Production 6 files to apply to the simulation [156]. When those become available the distributions should be re-examined with the weights applied. The largest backgrounds come from neutron interactions (specifically the photons and protons from neutron capture or scattering) and have their distribution highlighted in the plots. Truth information from the simulation was studied to quantify the total makeup of the backgrounds, and the results are shown for the ND in Table 6.1.

Particle Type	Source	% of Total Background
Photon	neutron capture	30.0
Photon	$\pi_0$ decay	18.2
Primary $\mu$	$\nu_\mu$ CC interaction	16.3
Proton	neutron elastic scattering	14.2
Proton or photon	neutron inelastic scattering	8.6
Primary proton	$\nu_\mu$ CC interaction	2.5

**Table 6.1:** Tabulation of the largest sources of background in reconstructed Michel electrons at the ND. Remaining percent is made up of charged particles and photons resulting from a large variety of different interactions, each comprising less than 1% of the total. The numbers are determined from GEANT4 truth information.

---

<sup>4</sup>Note that the theoretical true Michel distribution is not a Gaussian, and has a sharp cutoff at 53 MeV. The different shape we see results from imperfect energy reconstruction. For example, missing Michel energy due to traversing dead material. Essentially, our energy resolution for the Michels is not precise enough to resolve the true theoretical shape of the distribution.

### 6.4.3 Far Detector Michel electrons

Up until now only Michel electrons in the near detector had been examined in detail. I also took a first look at Michel electrons specifically in the FD. We expected to see differences at least in the timing distributions due to the coarser timing resolution and longer APD deadtime, and was curious whether they warranted creating separate MID templates for the FD. The energy,  $\Delta T$ , and MID distributions for the simulated FD Michel electrons are plotted in Fig. 6.36. There are several key differences to take note of. First, in the  $\Delta T$  distributions the effect of the timing resolution is evident. We see a small bump in events near the FD retrigger threshold of 4500 ns, and a more rounded peak for the events that don't retrigger, with the peak occurring at a later  $\Delta T$  than in the ND. In the energy distribution, we see the same spike at low-energy from the backgrounds, but we also see a smaller increase at low energy for the true Michel electrons. This is related to the timing differences as well. Because of the longer APD deadtime, many of the Michel electrons will appear while the cell that the muon decayed in is still dead. Any energy deposited by the Michel in that cell will then be missed if it occurs during the deadtime. Since the Michels only make a handful of hits to begin with, this can represent a significant percent of the total Michel energy. This effect becomes more obvious when we plot the energy of true Michels that retrigger vs. those that don't in the FD. This is shown in Fig. 6.37. The Michels that decayed late enough to register a retriggered hit have an energy distribution that more closely resembles the expected true Michel energy distribution, while the Michels that do not retrigger have the excess at lower energy from events that missed hits.

With significant differences in both the timing and energy distributions, we elected to make separate MID template histograms for the FD Michels. These are shown in Fig. 6.38. The true Michel templates have a similar shape to their ND counterparts, but there are more events concentrated at low NCells/CalE, and fewer events at low  $\Delta T$ . There are also more events concentrated in the higher DistToSlc bins because of the longer APD deadtime. The non-Michel background templates look the same as the ND for the NCell/CalE

Particle Type	Source	% of Total Background
Photon	neutron capture	88.8
Photon	$\pi_0$ decay	3.2
Primary $\mu$	$\nu_\mu$ CC interaction	2.9
Proton	neutron elastic scattering	1.1
Proton or photon	$\mu^-$ capture	1.3

**Table 6.2:** Tabulation of the largest sources of background in reconstructed Michel electrons at the FD. Remaining percent is made up of charged particles and photons resulting from a large variety of different interactions, each comprising less than 1% of the total. The numbers are determined from GEANT4 truth information.

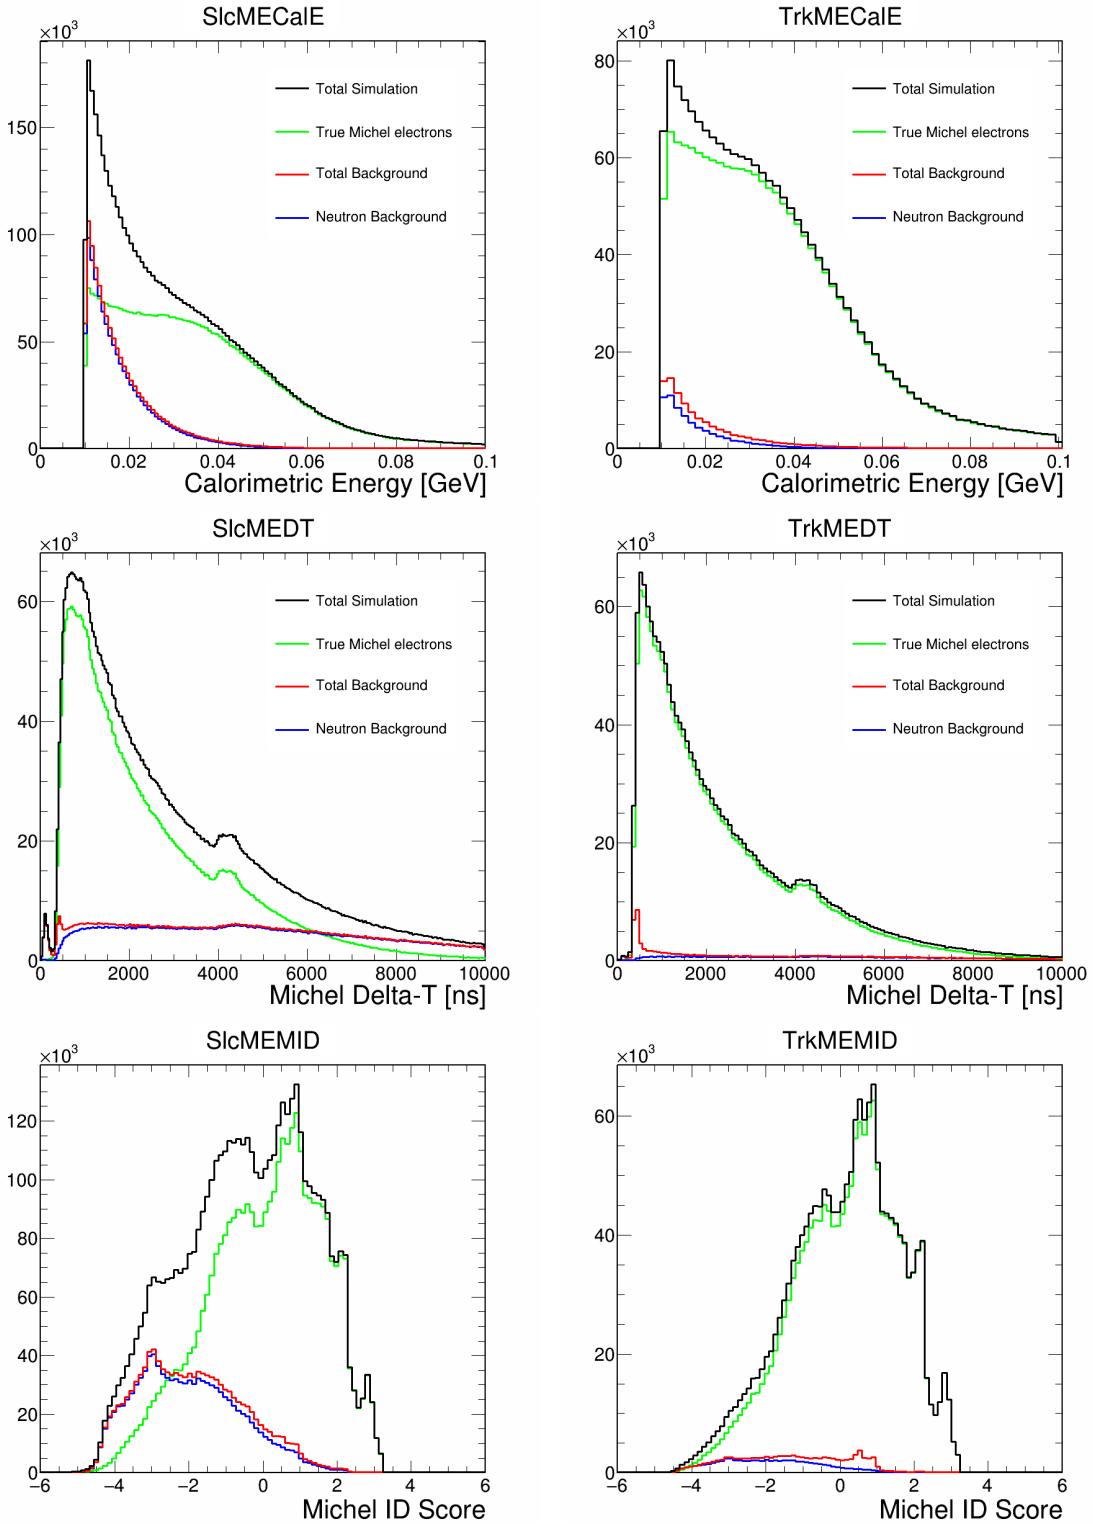
template, however the  $\text{DistToSlc}/\Delta T$  template shows a more diffuse distribution of events than the ND, where the backgrounds are concentrated at low  $\Delta T$ . This is a reflection of the different background composition at the FD. As can be seen in Fig. 6.36, almost all of the backgrounds at the FD are from neutron interactions. The specific backgrounds were again quantified from truth information and are summarized in Table 6.2. Similar types of background events are present but take up a smaller percent of the total than at the ND. Once again this can be related to the timing differences. As can be seen in the  $\Delta T$  plots for the ND (Fig. 6.35) there are large spikes in the non-neutron backgrounds at low  $\Delta T$ , which are significantly reduced at the FD. These additional backgrounds at very low  $\Delta T$  could be introduced from pile-up, which is not present at the FD.

As a final check on the impact of these improvements, we calculated the difference in Michel electron tagging efficiency before and after the changes to simulation/reconstruction. The tagging efficiency is defined as the sum of all reconstructed true Michel electrons (both SlcME and TrkME) divided by the total true Michel electrons that made hits in the detector, across all events. The result is shown in Table 6.3 for both ND and FD, before and after the improvements. The ND sees an increase from 71.7% to 84.8% efficiency, which corresponds to finding 18.3% more Michel electrons. The FD sees a modest increase from 59.3% to 61% efficiency when using the FD MID templates in MEFinder. Most of the improvement seems to come from accessing the additional phase space below 800 ns  $\Delta T$ . The FD, because of the longer APD deadtime, doesn't gain much from this, and has a

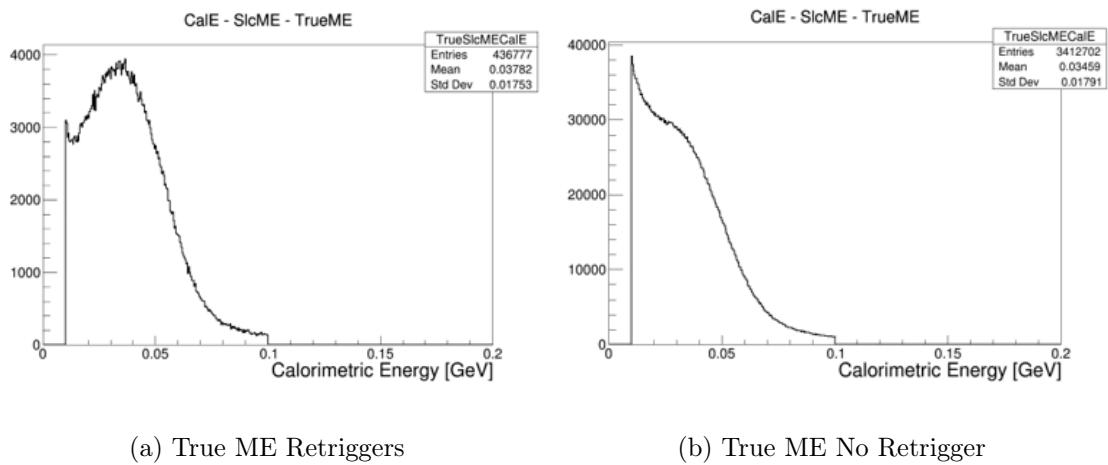
lower overall tagging efficiency as well compared to the ND. That being said, this number does not take into account any improvements to the data/MC agreement. While we didn't see a large increase in the tagging efficiency at the FD, we are no longer applying analysis cuts on  $\Delta T$  which will increase the amount of Michel electrons we can use for things like the Michel decomposition.

Detector	Tagging Efficiency (Old Sim/Reco)	Tagging Efficiency (Updated Sim/Reco)
ND	71.7	84.8
FD (using ND templates)	59.3	59.7
FD (using FD templates)		61.0

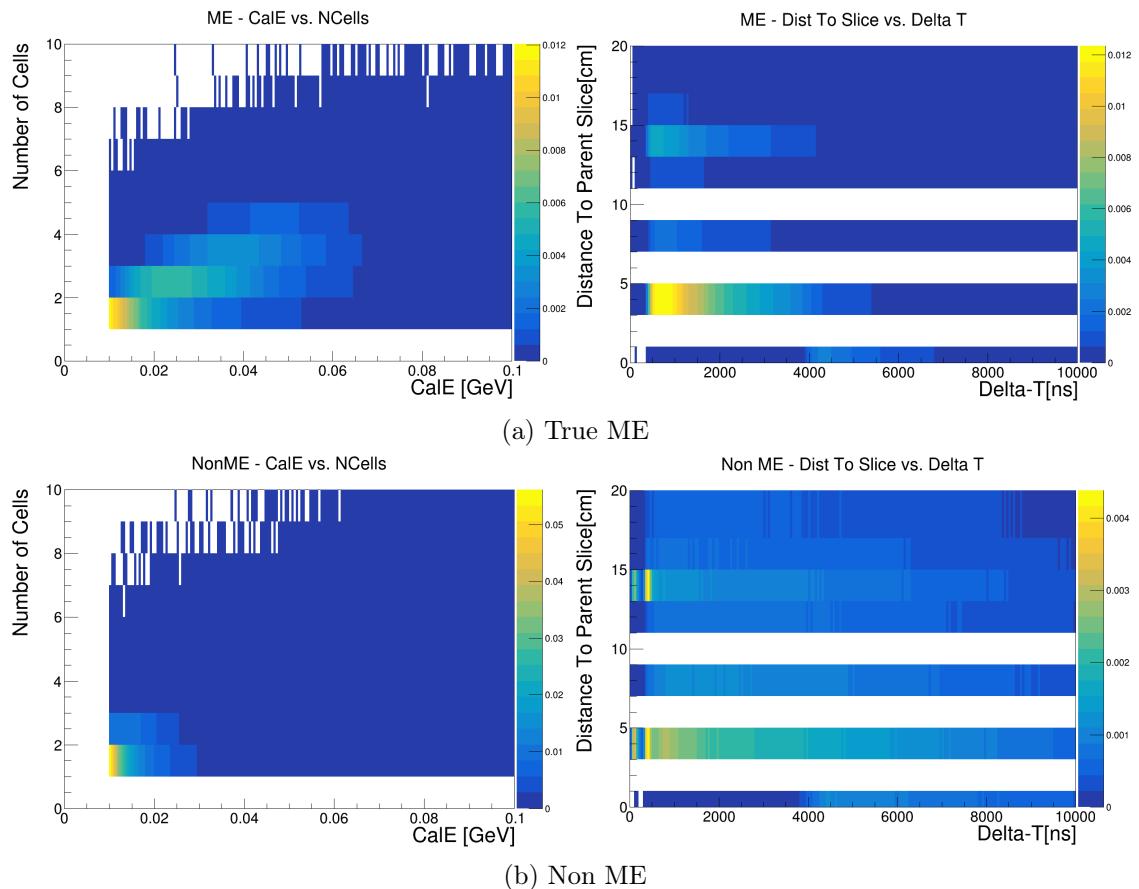
**Table 6.3:** Michel electron tagging efficiency before and after updates to improve the simulation and reconstruction of Michels. Shown separately for each detector, with the FD using either the ND MID templates for slice-matching, or the new FD-specific MID templates.



**Figure 6.36:** Reconstructed energy,  $\Delta T$ , and MID distributions for SlcME (left) and TrkME (right) in the far detector. These Michels were reconstructed using the ND MID templates in Fig. 6.34.



**Figure 6.37:** Reconstructed energy of true SlcME Michel electrons in the FD, split by those that retrigger (left) or don't retrigger (right) the parent slice. The spike at low energy for those that don't retrigger is from Michels that miss hits due to the longer APD deadtime in the FD.



**Figure 6.38:** FD MID template histograms made with all simulation and reconstruction updates applied. The top two plots are for true Michel, and the bottom two are for the non-Michel backgrounds.

#### 6.4.4 Future improvements using Michel electrons

Looking ahead, we are considering new ways to utilize Michel electrons in the 3-flavor analysis that can leverage my improvements. One idea is to use Michel electrons as a way to constrain wrong-sign backgrounds in the FD. Wrong-sign events are considered an irreducible background because we cannot distinguish the charge of the particles produced. However, because of the different decay behavior of muons/antimuons, we can expect different rates of Michel electrons between neutrino/antineutrino events. A  $\mu^+$  will always decay and produce a Michel electron, but a  $\mu^-$  will sometimes get captured on a nucleus instead of decaying, and not produce a Michel electron. Therefore, we should expect a higher rate of Michels from  $\bar{\nu}_\mu$  CC events (which produce  $\mu^+$ ) than  $\nu_\mu$  CC events (which produce  $\mu^-$ ). Michels can also appear in  $\nu_e$  ( $\bar{\nu}_e$ ) events through pion production.  $\nu_e$  CC events can produce  $\pi^+$  through DIS or RES interactions, while  $\bar{\nu}_e$  can produce  $\pi^-$ . The pions can decay, producing muons that can then decay to a Michel. However, the  $\pi^-$  and  $\mu^-$  can both be captured on a nucleus, which will result in no Michel, whereas  $\pi^+$  will be more likely to decay to a  $\mu^+$ , which will then always decay to a Michel. From this we would expect a higher rate of Michels from  $\nu_e$  CC events than from  $\bar{\nu}_e$  CC events.

To get an idea of the difference in the rates for NOvA, we studied the truth information of neutrino events in the Far detector. We looked at the number of true  $\nu_\mu$  ( $\bar{\nu}_\mu$ ) and  $\nu_e$  ( $\bar{\nu}_e$ ) events that passed the full 2024  $\nu_e$  and  $\nu_\mu$  selections in the FD, and counted how many of them contained true Michel electrons that made hits in the detector <sup>5</sup>. The calculation is done for both FHC and RHC beam modes. The numbers for  $\nu_\mu$  events are shown in Table 6.4, and for  $\nu_e$  events in Table 6.5. The  $\nu_\mu$  CC events have a higher rate of Michel electrons, with a small difference in the rates between  $\nu_\mu$  signal and wrong-sign (WS). The  $\nu_e$  CC events produce fewer Michels, but there is a more significant difference between the  $\nu_e$  signal/WS Michel rates than in the  $\nu_\mu$  CC events. A constraint would likely be most effective for the RHC  $\nu_e$  sample, where the WS makes up a larger portion of the sample,

---

<sup>5</sup>Note that this does not imply that the Michels were successfully reconstructed, just that they made hits in the detector, so there is a chance we could reconstruct them.

(a) FHC			
Channel	Total Events	Events w/ True Michel	% of Total
$\nu_\mu \rightarrow \nu_\mu$ (Sig)	316.63	285.05	90.0%
$\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$ (WS)	20.96	20.96	100%

(b) RHC			
Channel	Total Events	Events w/ True Michel	% of Total
$\nu_\mu \rightarrow \nu_\mu$ (WS)	22.11	20.12	91.0%
$\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$ (Sig)	64.03	64.03	100%

**Table 6.4:** The number of true signal and wrong-sign  $\nu_\mu$  events that contain a visible true Michel electron compared to all events. Events are passing the full  $\nu_\mu$  selection described in 5.2.3. Predictions were generated without extrapolation using the best-fit oscillation parameters found in Ch. 7.

(a) FHC			
Channel	Total Events	Events w/ True Michel	% of Total
$\nu_\mu \rightarrow \nu_e$ (Sig)	86.53	24.97	28.9%
$\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ (WS)	1.31	0.091	7.0%

(b) RHC			
Channel	Total Events	Events w/ True Michel	% of Total
$\nu_\mu \rightarrow \nu_e$ (WS)	1.62	0.465	28.7%
$\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ (Sig)	13.42	0.670	5.0%

**Table 6.5:** The number of true signal and wrong-sign  $\nu_e$  events that contain a visible true Michel electron compared to all events. Events are passing the full Core  $\nu_e$  selection described in 5.2.2. Predictions were generated without extrapolation using the best-fit oscillation parameters found in Ch. 7.

and there is the largest discrepancy between the Michel electron rates.

The exact nature of the constraint we can implement will depend on how much of this effect we can see in reconstructed Michel electrons. If a larger effect is seen, we could incorporate Michel electron info into a selection to create a new WS-enhanced subsample or reclaim events. The Michel info on its own is unlikely to be a sufficient selection but could be combined with some of the ideas discussed earlier in the chapter for reclaiming events. If a smaller effect is seen we could instead do a fit to the observed Michel rates while varying the amount of WS, to check that the amount of WS is within expected levels. If there appears to be a discrepancy in the amount of WS we are estimating then

we could include an additional systematic uncertainty. If not, this would remain as a simple cross-check on the amount of WS background.

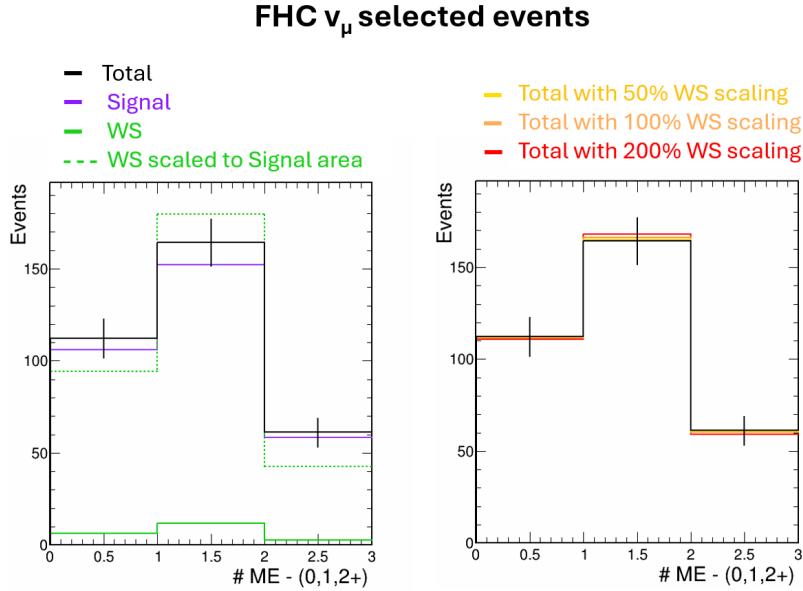
We then studied the number-of-Michel-electron (NMichel) distributions for reconstructed  $\nu_e$  and  $\nu_\mu$  events in the FD, for both FHC and RHC beam modes. The goal was to see how much WS was present in bins with reconstructed Michel electrons, relative to the signal, and see if it matched expectations from the truth study. Then, we can scale up the proportion of WS events in the sample, while keeping the same total, and see whether it results in a significant change to the observed NMichel distributions. How much the distribution changes indicates how significant of an effect a Michel constraint on the WS could have.

First we consider the FD  $\nu_\mu$  events. Fig. 6.39 shows the NMichel distribution for FD FHC simulation, for events passing the full  $\nu_\mu$  selection. From the truth study, we would expect the WS events in this sample to have a slightly higher rate of Michel electrons than the signal events. That should manifest as relatively more WS events in the 1 and 2+ NMichel bins. In the left hand plot, we see the total events, with signal and WS components indicated in purple and green. Since the WS is much smaller than the signal, in order to compare their shapes we also plot the WS scaled to match the area of the signal. Comparing the scaled WS and signal, we do see relatively more WS in the 1 Michel bin, matching the expectation from the truth study. However, in the 2+ Michel bin we see the opposite, with relatively less WS than signal. We have not yet determined a clear explanation, but one idea is that there could be more non-Michel backgrounds in that bin which behave differently from the true Michels.

Next we attempt to determine how much sensitivity the reconstructed Michel electrons have to the WS rate, by scaling up the proportion of WS events and seeing how the NMichel distributions change <sup>6</sup>. The right hand plot of Fig. 6.39 shows the effect of scaling up the proportion of WS events on the total. In order to see an effect, the WS needed to be scaled

---

<sup>6</sup>To give an example of how the scaling works, imagine there are 100 events total, 90 signal and 10 WS. If we scale up the WS 100% then we go from 10 → 20 WS and the signal scales down 90 → 80 to keep the same total events in the sample. We only considered an increasing WS proportion for this study.

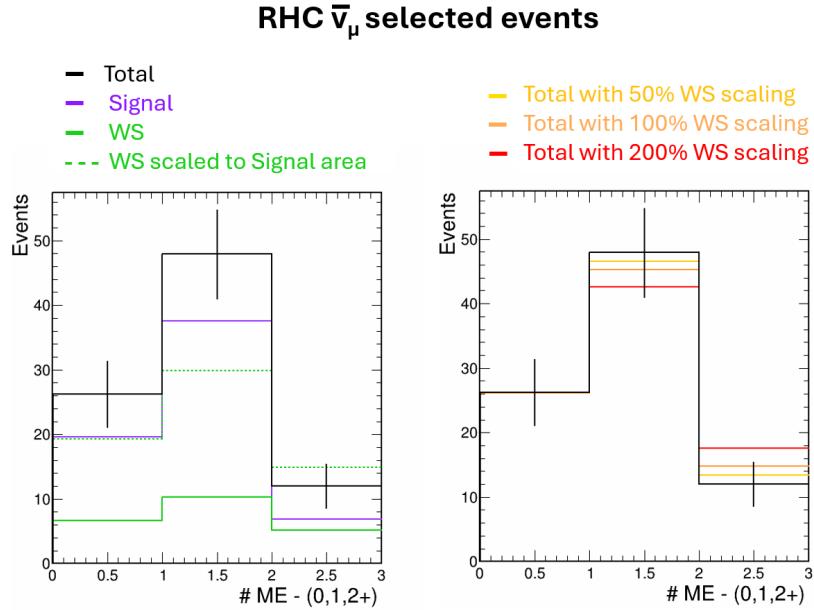


**Figure 6.39:** Simulated FD FHC events with 0, 1, or 2+ reconstructed Michel electrons, for events passing the full FD  $\nu_\mu$  selection. All histograms are scaled to this analysis' POT. The total prediction is shown in black with statistical error bars ( $\sqrt{N}$ ), and is the same in both histograms. The left shows the signal ( $\nu_\mu$ ) and WS ( $\bar{\nu}_\mu$ ) components of the total, as well as what the WS distribution would look like scaled up to match the size of the signal. The right shows what the total distribution would look like if we scaled up the WS by 50%, 100%, and 200% while keeping the same total number of events.

up by a large amount: 50%, 100%, and 200% <sup>7</sup> indicating that the Michels in this sample may not have much sensitivity. With 200% scaling we see a small increase in the 1 Michel bin, and a small decrease in the 0 and 2+ bins.

The corresponding distributions for the RHC  $\nu_\mu$  sample are shown in Fig. 6.40. From the truth study, we expect the signal events in RHC to have relatively more Michel electrons. This would be seen as relatively fewer WS in the 1 and 2+ NMichel bins. Comparing the scaled WS to the signal, this expectation is true for the 1 Michel bin, but reversed for the 2+ Michel bin. The RHC samples have a larger proportion of WS events, so we see more of an effect when scaling the WS. In the right-hand plot, the 50% scaling now has a small but visible effect, and the 200% scaling shows a large effect.

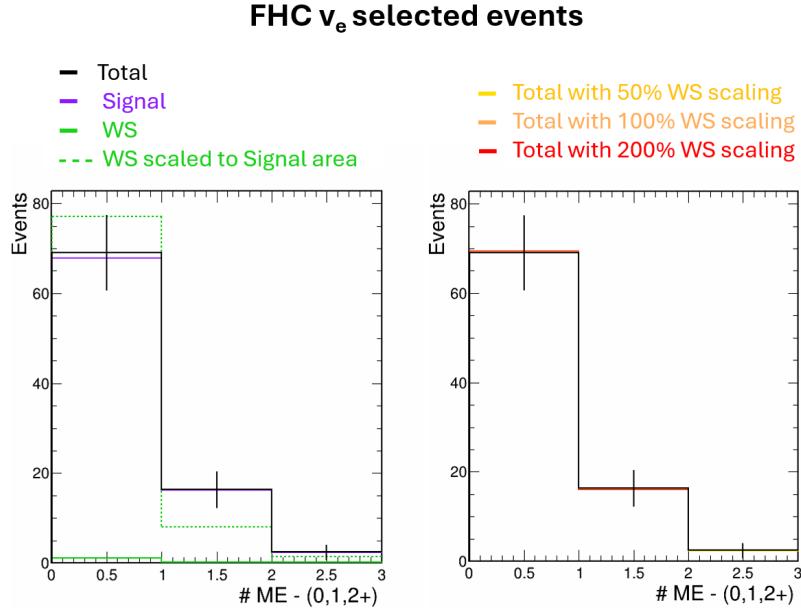
<sup>7</sup>We do not expect our WS to be off by this much. The large deviations are chosen only to understand the limit of our sensitivity using the Michels.



**Figure 6.40:** Simulated FD RHC events with 0, 1, or 2+ reconstructed Michel electrons, for events passing the full FD  $\bar{\nu}_\mu$  selection. All histograms are scaled to this analysis' POT. The total prediction is shown in black with statistical error bars ( $\sqrt{N}$ ) and is the same in both histograms. The left shows the signal ( $\bar{\nu}_\mu$ ) and WS ( $\nu_\mu$ ) components of the total, as well as what the WS distribution would look like scaled up to match the size of the signal. The right shows what the total distribution would look like if we scaled up the WS by 50%, 100%, and 200% while keeping the same total number of events.

Next we consider the  $\nu_e$  samples. Fig. 6.41 shows the NMichel distributions for FD FHC simulation, for events passing the full FD Core  $\nu_e$  selection. From the truth study, we expect the WS to have relatively fewer Michels than the signal, however we expect overall very few WS events. The shape of the scaled WS histogram does match the expectation from the truth study, but there are so few WS events that even scaling by 200% has a negligible effect on the total.

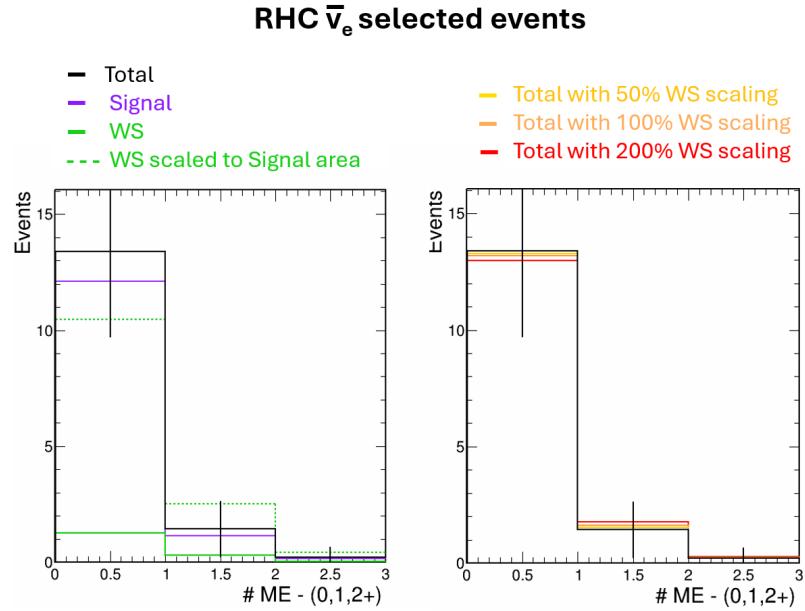
Fig. 6.42 shows the corresponding RHC  $\nu_e$  distributions. From the truth study we expect the WS to have relatively more Michels. We also expected this sample to have the most potential for sensitivity to the WS since there is both a larger amount of WS events, and the difference in Michel rate between signal/WS is the highest. In the left-hand plot, the signal and WS are actually at a comparable level even before scaling the WS. We see



**Figure 6.41:** Simulated FD FHC events with 0, 1, or 2+ reconstructed Michel electrons, for events passing the full FD Core  $\nu_e$  selection. All histograms are scaled to this analysis' POT. The total prediction is shown in black with statistical error bars ( $\sqrt{N}$ ) and is the same in both histograms. The left shows the signal ( $\nu_e$ ) and WS ( $\bar{\nu}_e$ ) components of the total, as well as what the WS distribution would look like scaled up to match the size of the signal. The right shows what the total distribution would look like if we scaled the WS up by 50%, 100%, and 200% while keeping the same total number of events.

more WS in both 1 and 2+ Michel bins, matching the expectation from the truth study. Despite this, we still require a large scaling of the WS to see a significant difference in the total NMichel distributions.

The reconstructed NMichel distributions show varying agreement with the expectations from the truth study. All the distributions agree with expectation in the 1 NMichel bin, but some differ in the 2+ bin for an unknown reason. This should be explored further to determine the cause. In terms of the effectiveness of using the Michel rates for a constraint on the WS, it seems as though the WS would need to be off by a significant amount, much more than we expect, for there to be a visible effect in the NMichel distributions. So there would not be much power in constraining the WS this way. The distributions that have been shown are for events passing selections, so there may still be some potential for the



**Figure 6.42:** Simulated FD RHC events with 0, 1, or 2+ reconstructed Michel electrons, for events passing the full FD Core  $\bar{\nu}_e$  selection. All histograms are scaled to this analysis' POT. The total prediction is shown in black with statistical error bars ( $\sqrt{N}$ ) and is the same in both histograms. The left shows the signal ( $\bar{\nu}_e$ ) and WS ( $\nu_e$ ) components of the total, as well as what the WS distribution would look like scaled up to match the size of the signal. The right shows what the total distribution would look like if we scaled the WS up by 50%, 100%, and 200% while keeping the same total number of events.

Michel electrons to be used to reclaim events. The Michel distributions for those events failing selections should be studied further.

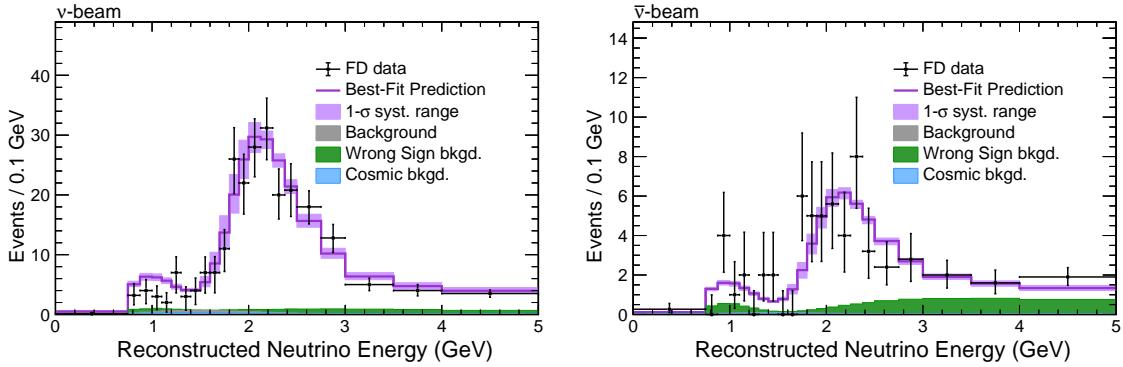
# Chapter 7

## 3-Flavor Analysis Results

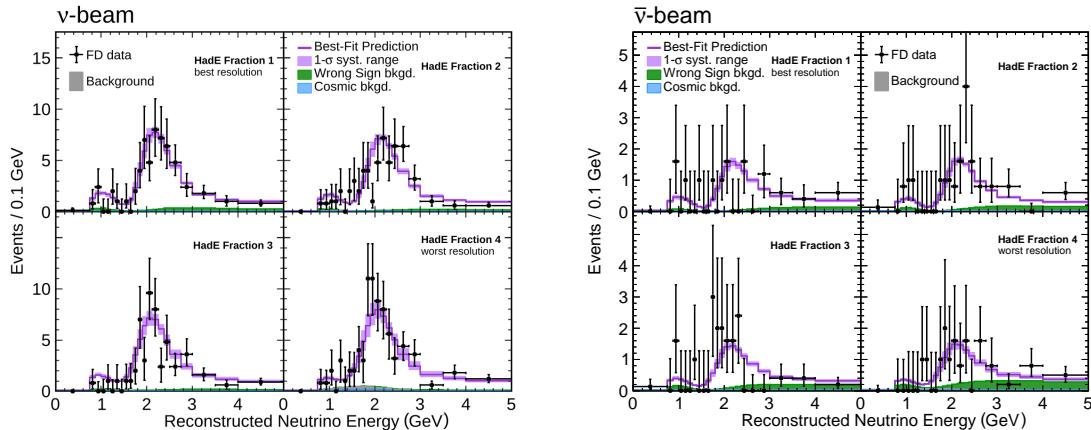
In this chapter we present the results of the joint analysis of  $26.61 \times 10^{20}$  POT of neutrino mode, and  $12.5 \times 10^{20}$  POT of antineutrino mode beam data. As outlined in Chapter 5, the predicted FD  $\nu_\mu$  and  $\nu_e$  energy spectra, along with the full list of systematic uncertainties, are fit at once to the observed FD data spectra using either frequentist or Bayesian techniques, in order to extract a measurement of the oscillation parameters. The  $\nu_\mu$  disappearance channel primarily gives us sensitivity to  $\Delta m_{32}^2$  and  $\sin^2 \theta_{23}$ , while the  $\nu_e$  appearance channel gives us sensitivity to the mass ordering, octant of  $\theta_{23}$ , and value of  $\delta_{\text{CP}}$ . We start by showing the observed data in the FD. We then present the results of the frequentist fit to the data. At this time Feldman-Cousins corrections are only available for a handful of our measurements, due to the significant computational cost of producing them. The majority of the frequentist results are shown with confidence levels drawn assuming the Gaussian approximation. While the contours may not exactly contain the stated coverage, they can still serve as a useful approximation of our results, and be used to make comparisons between contours under different fitting conditions. The remaining sections incorporate both frequentist and Bayesian results. We show the impact of the different Daya Bay reactor neutrino constraints on our measurements, and finally, conclude with a discussion of the new results in context with other results from the field of neutrino oscillation physics.

## 7.1 The Observed Data

Using the  $\nu_\mu$  selections outlined in Sec. 5.2.3, we observe 384 events in the data, with a predicted best-fit total of 408.6 (including 11.0 background). We observe 106  $\bar{\nu}_\mu$  events, with a predicted best-fit total of 97.7 (including 1.7 background). The energy spectra for the events are shown in Fig. 7.1, with all hadronic energy bins combined. Figure 7.2 shows the same distributions split into their hadronic energy bins. In the absence of neutrino

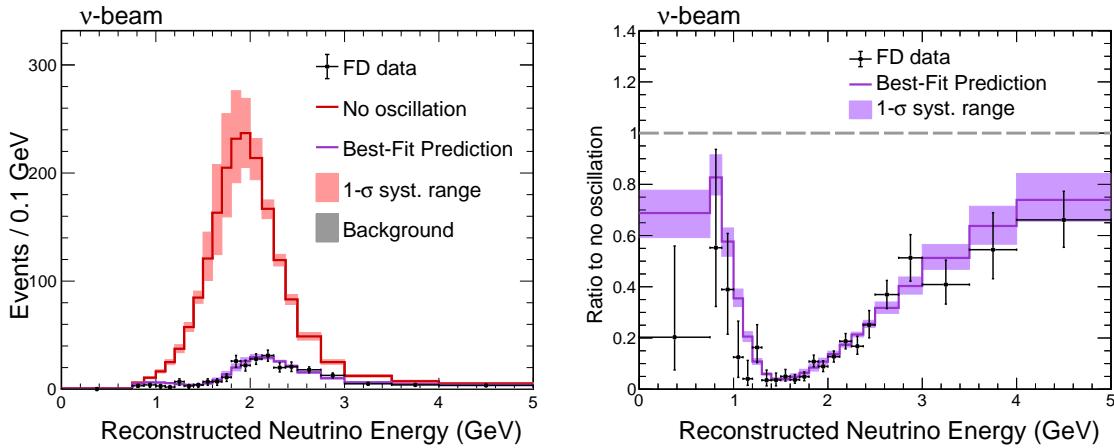


**Figure 7.1:** Observed and predicted energy spectra, for  $\nu_\mu$  (left) and  $\bar{\nu}_\mu$  (right) CC selected events in the FD. All hadronic energy fraction and transverse momentum quantiles have been combined in these spectra. The prediction is generated at the frequentist best-fit point.



**Figure 7.2:** Observed and predicted energy spectra, for  $\nu_\mu$  (left) and  $\bar{\nu}_\mu$  (right) CC selected events in the FD, split into four hadronic energy fractions but with the transverse momentum quantiles combined. The prediction is generated at the frequentist best-fit point.

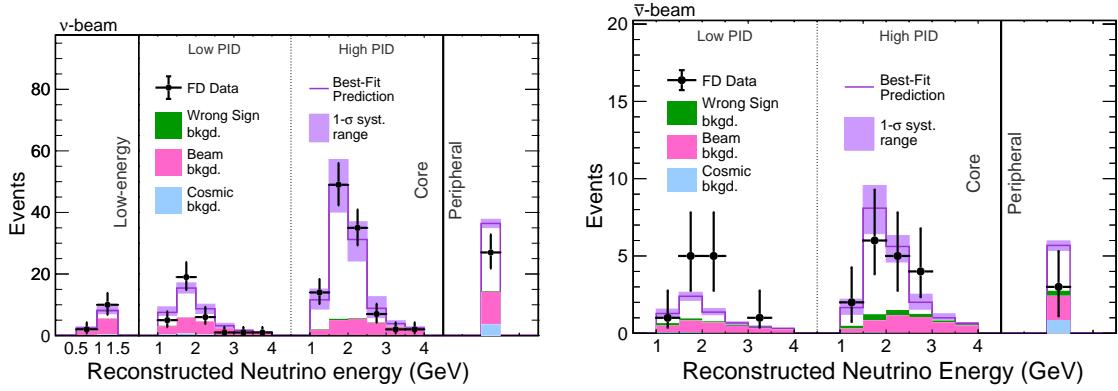
oscillations we would have expected to see  $\sim 2100$   $\nu_\mu$  events and  $\sim 500$   $\bar{\nu}_\mu$  events, so the observed event counts show a clear sign of muon neutrino disappearance in our beam. This can be visualized in Fig. 7.3 which plots the predicted FD energy spectra in the absence of oscillations on top of the observed spectra for the FHC beam. From this we can also plot the ratio of oscillated to unoscillated predictions, where we can clearly see the characteristic dip shape in the data.



**Figure 7.3:** **Left:** Observed and predicted energy spectra, for  $\nu_\mu$  CC selected events in the FD, for the oscillated (purple) and unoscillated (red) predictions. **Right:** Ratio of the observed data and oscillated prediction to the unoscillated prediction.

Using the  $\nu_e$  selections from 5.2.2 we observe 181 events in the data, with a predicted total of 191.2 events (including 62.5 background). We observe 32  $\bar{\nu}_e$  events, with a predicted total of 31.1 (including 12.2 background). The observation exceeds the expected background, giving a clear signal of  $\nu_e$  appearance in our beam. The  $\nu_e$  events for each of the selected samples are shown in Fig. 7.4. The total observed and predicted  $\nu_\mu$  ( $\bar{\nu}_\mu$ ) and  $\nu_e$  ( $\bar{\nu}_e$ ) event counts, along with the predicted background components, are summarized in Table 7.1.

Figures 7.5 and 7.6 show comparisons of the observed number of appearance events to the predicted number of events under different combinations of the mass ordering,  $\delta_{\text{CP}}$  and  $\sin^2 \theta_{23}$ . This can help give context to the full fit results in the next section. Figure 7.5 shows the FHC and RHC appearance samples separately. Neither FHC or RHC data



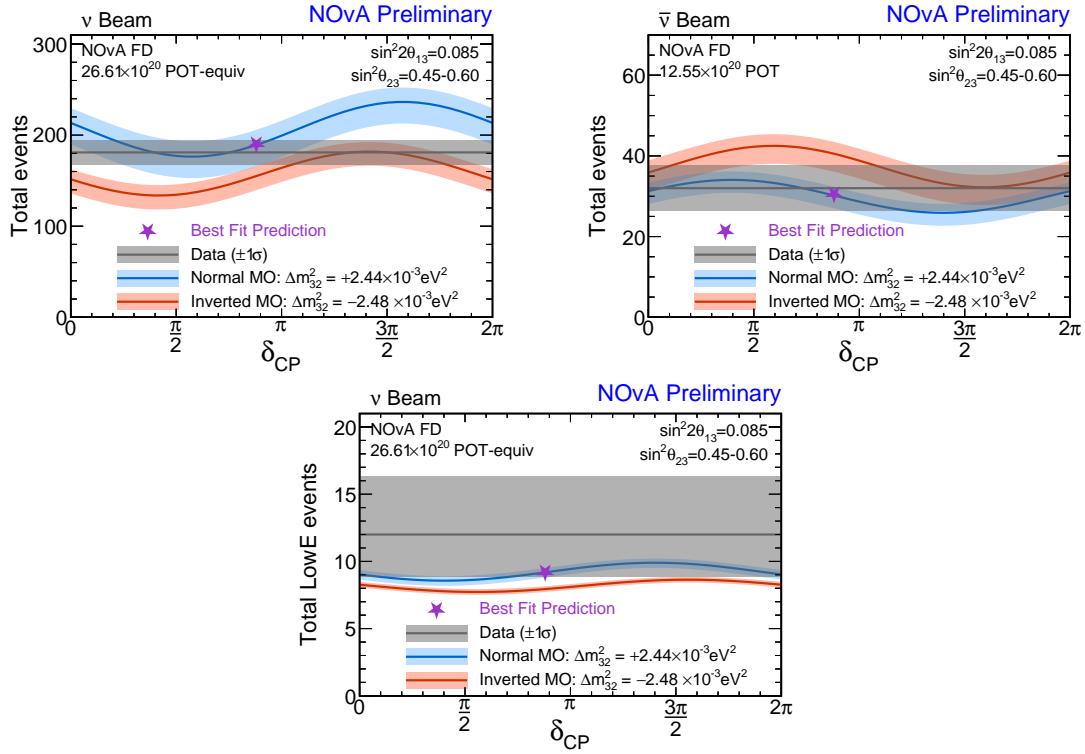
**Figure 7.4:** Observed and predicted energy spectra, for  $\nu_e$  (left) and  $\bar{\nu}_e$  (right) CC selected events in the FD. The prediction is generated at the frequentist best-fit point.

show a strong preference towards one of the mass orderings, being compatible with either one over different ranges of  $\delta_{CP}$ . The new low-energy  $\nu_e$  sample is also shown on its own. In this subsample we see slightly more data events than expected, however the best-fit prediction still falls within  $1\sigma$  of the data. The excess does give a slight preference towards the normal mass ordering but the effect on the total  $\nu_e$  sample is small.

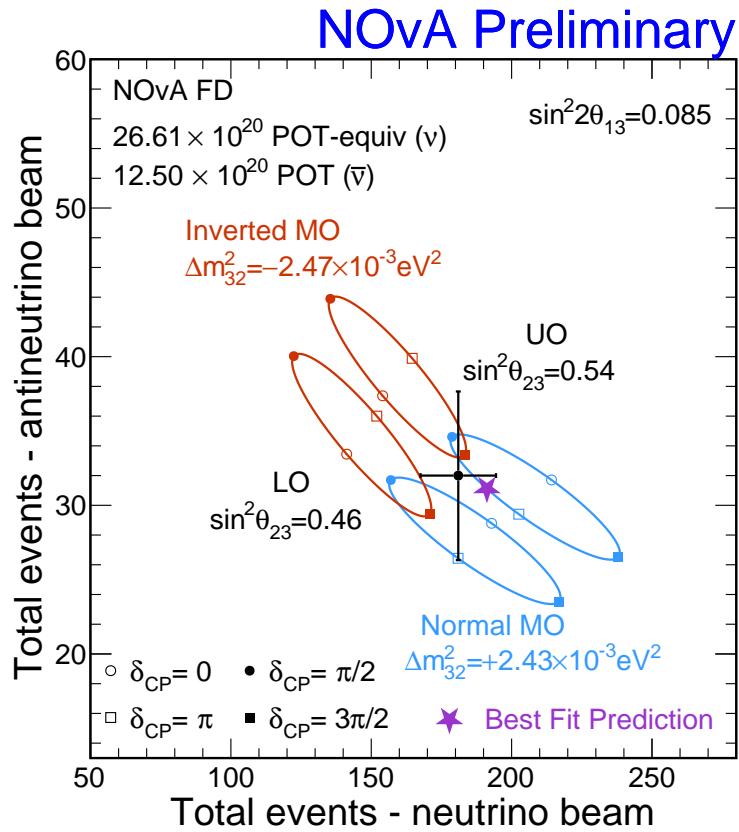
We can combine the FHC and RHC into a 2D representation (the bi-event plot) shown in Fig. 7.6. This shows the number of RHC appearance events on the y-axis versus the number of FHC appearance events on the x-axis. This allows us to look for asymmetry in the rates of  $\nu_e$  and  $\bar{\nu}_e$  appearance. The ellipses show the predicted event counts in each beam mode under different combinations of the oscillation parameters. With 181 FHC events and 32 RHC events, our data falls right in the middle region where the matter and CP violation effects are degenerate. In other words, we prefer combinations of the oscillation parameters that lead to similar probability of  $\nu_e$  and  $\bar{\nu}_e$  appearance, such as the normal mass ordering with  $\delta_{CP}$  near  $\pi/2$ , or the inverted mass ordering with  $\delta_{CP}$  near  $3\pi/2$ . We disfavor the combinations of oscillation parameters that would lead to a large asymmetry in the rates of  $\nu_e$  vs.  $\bar{\nu}_e$ , such as the normal mass ordering with  $\delta_{CP}$  near  $3\pi/2$ , or the inverted mass ordering with  $\delta_{CP}$  near  $\pi/2$ .

Sample	FHC			RHC	
	$\nu_\mu$	$\nu_e$	Low-energy	$\bar{\nu}_\mu$	$\bar{\nu}_e$
$\nu_\mu \rightarrow \nu_\mu$	372.3	4.3	0.3	24.4	0.2
$\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$	24.5	0.1	0.0	71.5	0.2
$\nu_\mu \rightarrow \nu_e$	0.4	125.3	3.4	0.0	2.1
$\bar{\nu}_\mu \rightarrow \bar{\nu}_e$	0.0	1.8	0.1	0.0	18.9
Beam $\nu_e + \bar{\nu}_e$	0.1	26.1	0.8	0.0	6.5
NC	5.5	16.8	5.3	0.8	2.0
Cosmic	4.4	5.5	0.5	0.7	1.1
Others	1.5	0.8	0.1	0.2	0.1
Signal	397.6	125.3	3.4	96.0	18.9
Background	11.0	55.4	7.1	1.7	12.2
Predicted	408.6	180.7	10.5	97.7	31.1
Observed	384	169	12	106	32

**Table 7.1:** Predicted event counts for the  $\nu_\mu$ ,  $\nu_e$  and Low-energy selected samples in the neutrino beam, and  $\bar{\nu}_\mu$  and  $\bar{\nu}_e$  samples in the antineutrino beam. The low/high PID, and Peripheral samples are combined in the  $\nu_e$  ( $\bar{\nu}_e$ ) columns, while the Low-energy sample is shown separately due to its novel status. Any oscillation channels not listed are in “Others”. Signal in the  $\nu_\mu$  ( $\bar{\nu}_\mu$ ) columns includes wrong-sign events and some  $\nu_\mu$  from Others. Predictions were generated using the best-fit oscillation parameters found from a frequentist fit to the data with the 1D Daya Bay constraint on  $\sin^2 2\theta_{13}$  [139].



**Figure 7.5:** Comparison of the predicted number of FD  $\nu_e$  and  $\bar{\nu}_e$  appearance events as a function of  $\delta_{CP}$  in each mass ordering. The solid red and blue lines represent the best-fit value of  $\sin^2 \theta_{23}$ , while the band shows the variation over a wider range of  $\sin^2 \theta_{23}$  values listed in the top right. The total observed data events are shown in gray along with the statistical  $1\sigma$  error bands. Shown for the full FHC (including LowE) and RHC samples on top, with the LowE sample additionally shown on its own below.

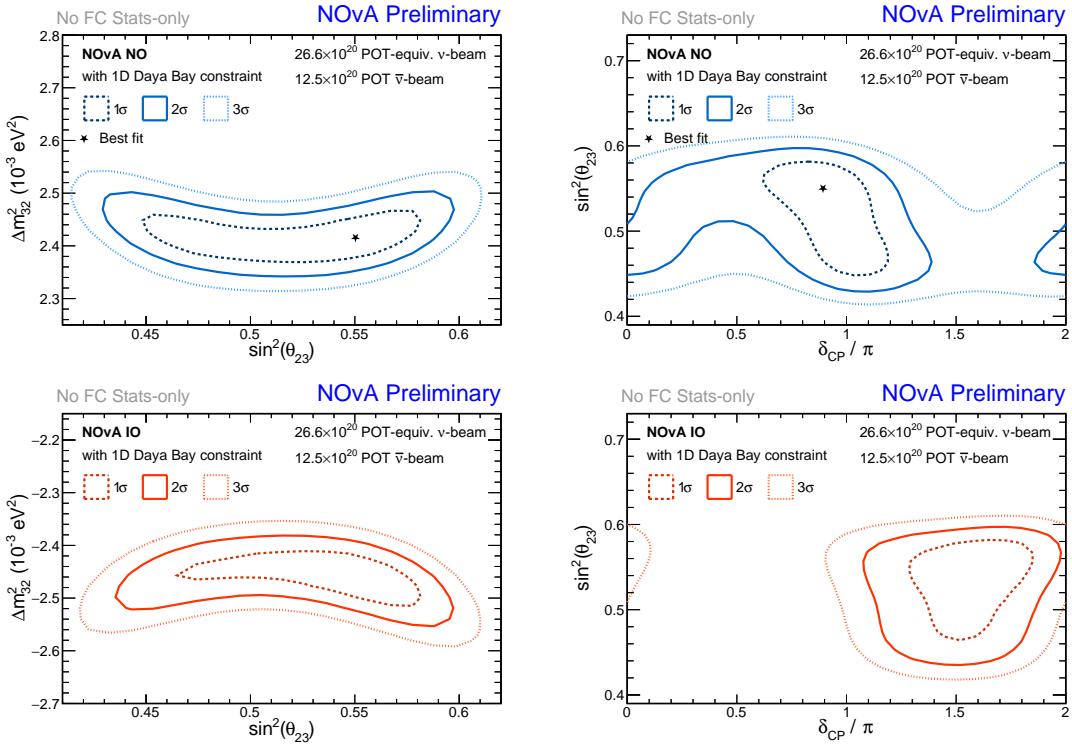


**Figure 7.6:** Bi-event plot showing total predicted appearance events in RHC and FHC for different combinations of octant and hierarchy. Each ellipse uses the values of mass ordering and octant indicated by the color and closest label. The value of  $\delta_{CP}$  is varied to trace out the ellipse, with important values marked. Observed number of events is shown with errors. The best-fit point in the UO, Normal MO is indicated. Systematic pulls are included in the predictions. Low-energy sample is included in both observed and predicted event counts. Our data falls in the middle region, away from combinations of oscillation parameters that lead to a large asymmetry in  $\bar{\nu}_e$  vs.  $\nu_e$  events.

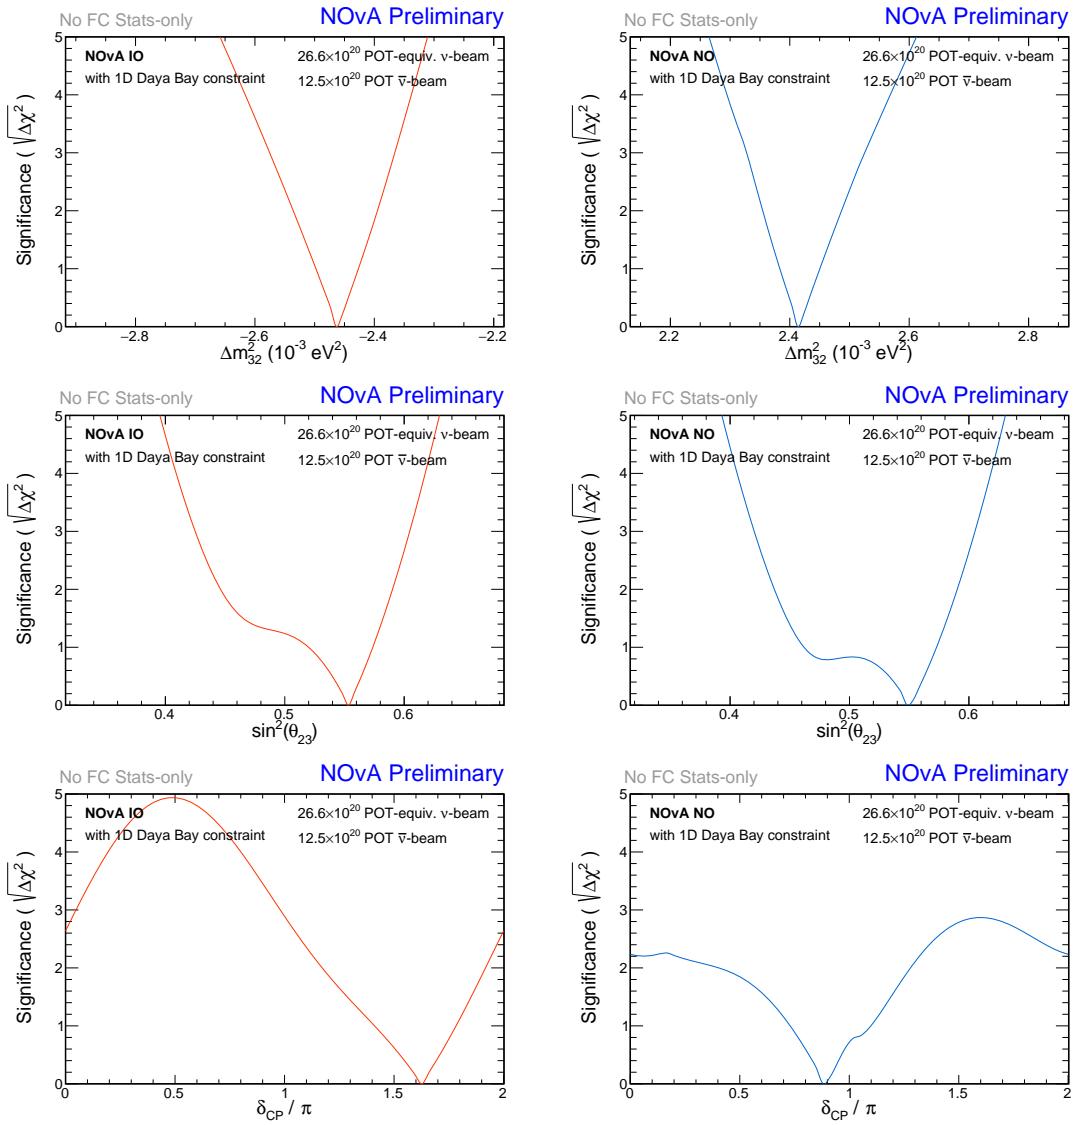
## 7.2 Fits and Results

We begin by presenting analysis results using the frequentist approach to fitting. First, doing a stats-only fit, and then incorporating all systematics. The available Feldman-Cousins corrections are shown in the subsequent section.

Figure 7.7 shows the 2D confidence-level contours in  $\Delta m_{32}^2$  vs.  $\sin^2 \theta_{23}$ , and  $\sin^2 \theta_{23}$  vs.  $\delta_{\text{CP}}$  space, from a frequentist fit to the data using the 1D Daya Bay constraint on  $\sin^2 2\theta_{13}$ . Figure 7.8 shows 1D fits for  $\Delta m_{32}^2$ ,  $\sin^2 \theta_{23}$ , and  $\delta_{\text{CP}}$ . The oscillation parameters not present on the axes of a given plot have been profiled over at each point in the space. These are stats-only fits, so the systematic parameters are not included in the log-likelihood minimization. We find a best fit point in the normal mass ordering and upper octant of  $\theta_{23}$ . The allowed values of  $\theta_{23}$  are compatible with maximal mixing in both mass orderings. As hinted at by Fig. 7.6, we disfavor regions of parameter space that lead to high asymmetry in  $\nu_e$  ( $\bar{\nu}_e$ ) appearance rates. Notably, in the inverted mass ordering we exclude most values of  $\delta_{\text{CP}}$  between 0 and  $\pi$  at greater than  $3\sigma$ .



**Figure 7.7:** Frequentist 1,2,3 $\sigma$  confidence level contours from stats-only fits to the full FHC and RHC datasets, using the 1D Daya Bay constraint on  $\sin^2 2\theta_{13}$ . Top panels show normal mass ordering (NO) in blue, and bottom panels show inverted mass ordering (IO) in red. The best-fit point is found in the normal mass ordering, and upper octant of  $\theta_{23}$ .

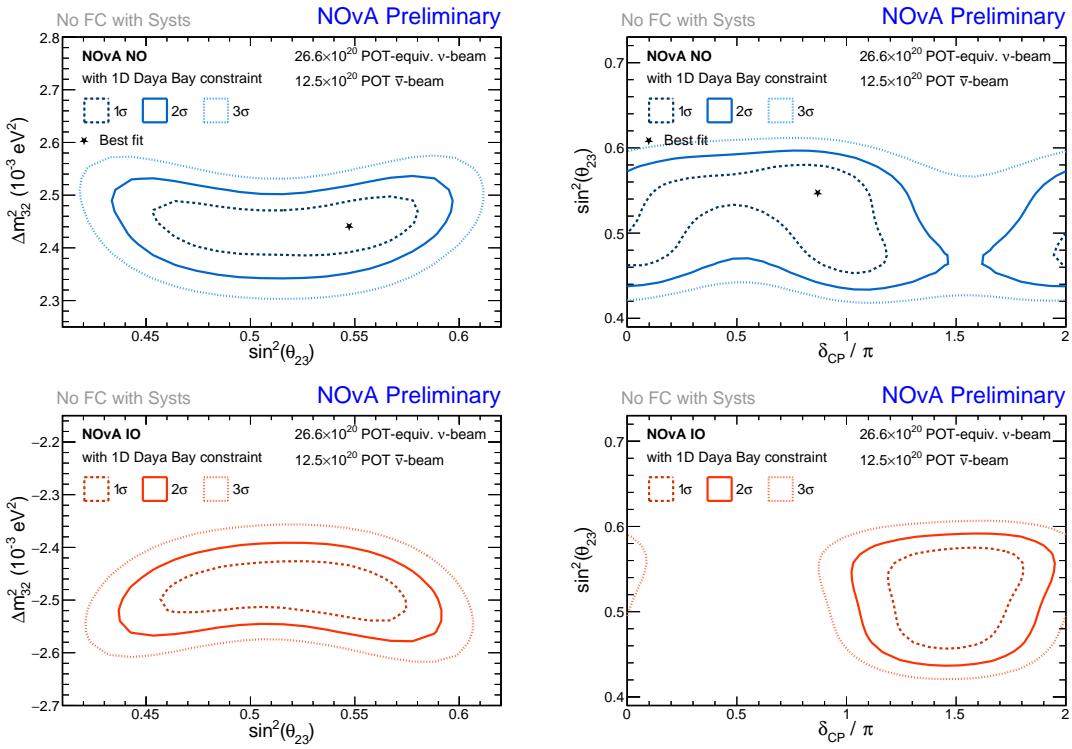


**Figure 7.8:** 1D stats-only frequentist fits to the full FHC and RHC datasets, using the 1D Daya Bay constraint on  $\sin^2 2\theta_{13}$ . Assuming the inverted mass ordering (left) and normal mass ordering (right).

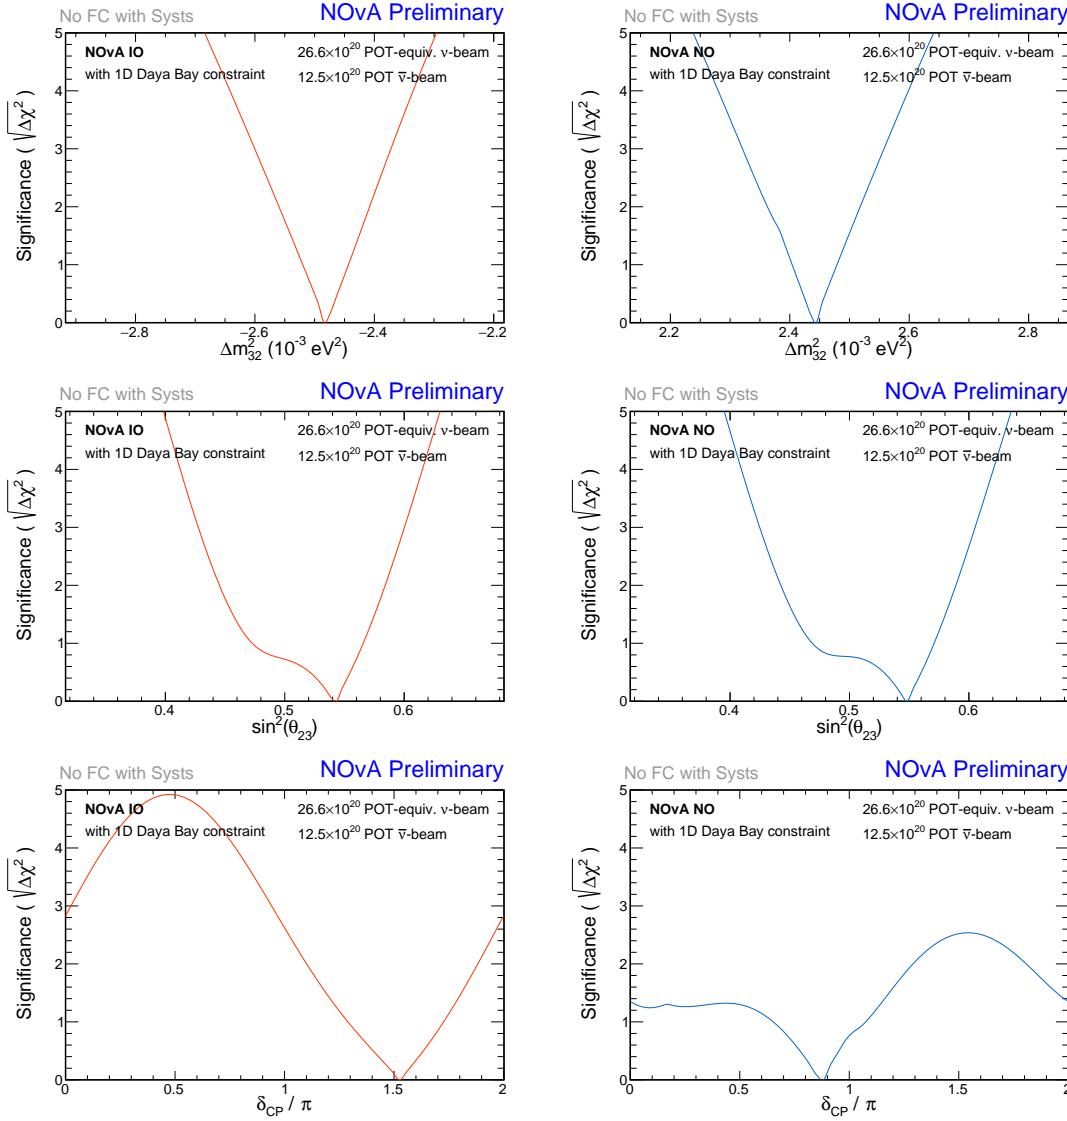
Next we repeat the fits, this time including the systematic terms in the log-likelihood minimization. At each point in the phase space of a given plot, we profile over all the systematic parameters as well as the other oscillation parameters. Figure 7.9 shows the resulting 2D confidence-level contours, and Fig. 7.10 shows the results of the 1D fits. These are the results on which Feldman-Cousins corrections will be applied. The inclusion of systematics broadens the contours slightly, while keeping the same general conclusions from the stats-only fits. With the systematics included in the fit, and using the 1D Daya Bay constraint on  $\sin^2 2\theta_{13}$ , we find a best-fit point in the normal mass ordering and upper octant of  $\theta_{23}$ , with the oscillation parameters

- $\Delta m_{32}^2 = 2.433$
- $\sin^2 \theta_{23} = 0.546$
- $\delta_{\text{CP}} = 0.87^\circ$

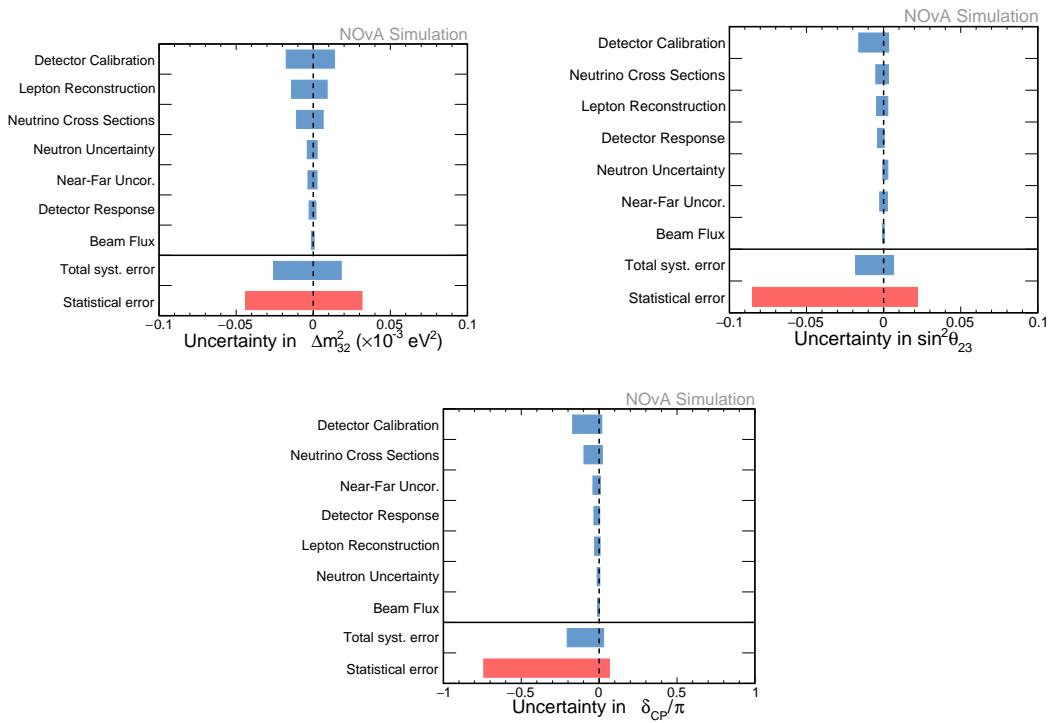
The systematic uncertainties associated with each of these parameters at the best-fit point is summarized in Fig. 7.11. Our largest systematic uncertainty comes from detector calibrations, followed by either lepton reconstruction (for the  $\Delta m_{32}^2$  measurement) or neutrino cross sections (for  $\delta_{\text{CP}}$  and  $\sin^2 \theta_{23}$ ). The statistical uncertainty is still dominant for all measurements.



**Figure 7.9:** Frequentist 1,2,3 $\sigma$  confidence level contours from fits to the full FHC and RHC datasets, profiling over all systematic uncertainties, and using the 1D Daya Bay constraint on  $\sin^2 2\theta_{13}$ . Top panels show normal mass ordering (NO) in blue, and bottom panels show inverted mass ordering (IO) in red.

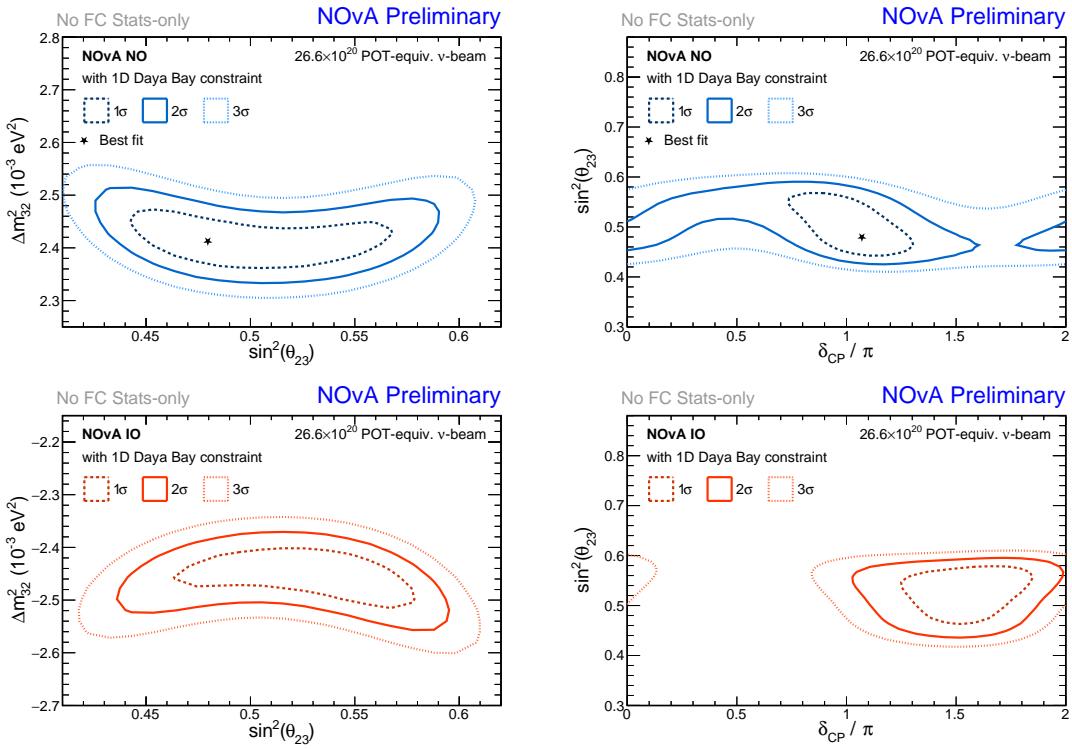


**Figure 7.10:** 1D frequentist fits to the full FHC and RHC datasets, profiling over all systematic uncertainties, and using the 1D Daya Bay constraint on  $\sin^2 2\theta_{13}$ . Assuming the inverted mass ordering (left) and normal mass ordering (right).

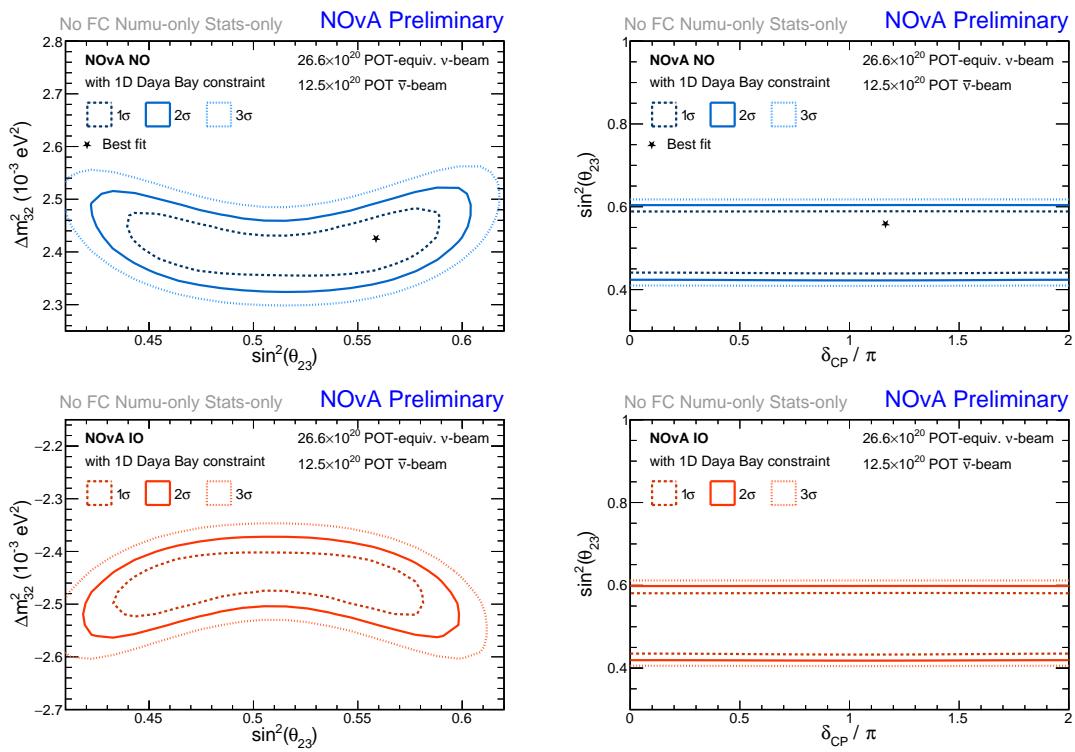


**Figure 7.11:** Systematic and statistical uncertainties on the measured values of  $\Delta m_{32}^2$ ,  $\sin^2 \theta_{23}$ , and  $\delta_{CP}$  at the best-fit point. From [157].

As an additional exercise, we can perform the fits while excluding a subset of our data (and corresponding predictions) and compare with the full result above. This can give us a better sense of the effect the excluded sample has on the full result. Two studies of interest are the results without the RHC data (an FHC-only fit) and the results without the  $\nu_e$  data (a  $\nu_\mu$ -only fit). The plots shown are for stats-only fits with the 1D Daya Bay constraint. Figure 7.12 shows the FHC-only fits. Compared to the full stats-only fit, the contours are slightly expanded and shifted down to lower values of  $\sin^2 \theta_{23}$ . Notably, the best-fit point has also changed from the upper to lower octant of  $\theta_{23}$ . Figure 7.13 shows the  $\nu_\mu$ -only fits. These include both FHC and RHC  $\nu_\mu$  ( $\bar{\nu}_\mu$ ) samples, but no  $\nu_e$  samples. From this we can clearly see how all of our sensitivity to  $\delta_{\text{CP}}$  comes from the  $\nu_e$  appearance measurement, since the contour is completely flat in that space now. For the  $\Delta m_{32}^2$  vs.  $\sin^2 \theta_{23}$  space, the contours have expanded slightly, and the best fit has shifted to a higher value of  $\sin^2 \theta_{23}$ .



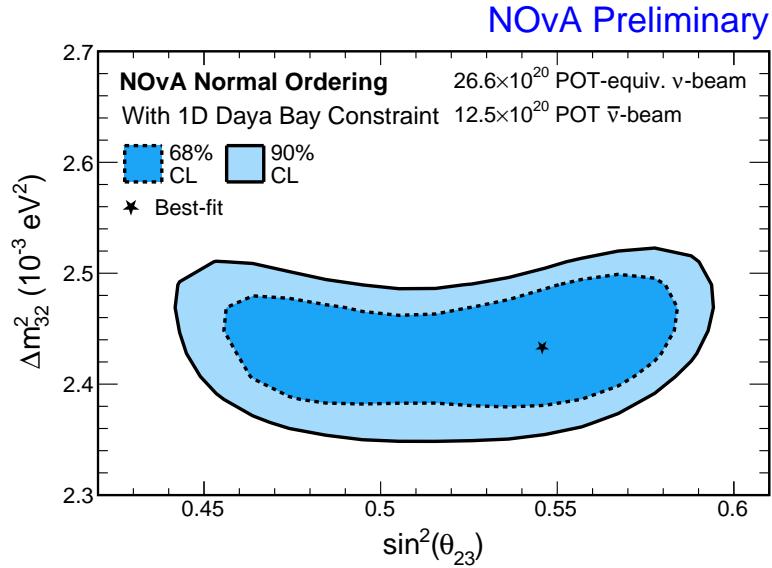
**Figure 7.12:** Frequentist 1,2,3 $\sigma$  confidence level contours from stats-only fits to the FHC-only  $\nu_\mu + \nu_e$  dataset, using the 1D Daya Bay constraint on  $\sin^2 2\theta_{13}$ . Top panels show normal mass ordering (NO) in blue, and bottom panels show inverted mass ordering (IO) in red.



**Figure 7.13:** 1,2,3 $\sigma$  confidence level contours from stats-only fits to the  $\nu_\mu$ -only dataset, using the 1D Daya Bay constraint on  $\sin^2 2\theta_{13}$ . Top panels show normal mass ordering (NO) in blue, and bottom panels show inverted mass ordering (IO) in red.

### 7.2.1 Feldman-Cousins Corrections

The frequentist results shown so far will need to be Feldman-Cousins corrected to ensure accurate coverage before final publication. Currently we have corrections for the  $\Delta m_{32}^2$  vs.  $\sin^2 \theta_{23}$  fit in the normal mass ordering, and a calculation of our preference for the normal mass ordering under different reactor constraints. Efforts are underway to produce corrections for the remaining measurements. Figure 7.14 shows the corrected surface. The best-fit oscillation parameters with corrected  $1\sigma$  uncertainties are shown in the table in Fig. 7.15, along with the significance with which we prefer the normal mass ordering under the 1D or 2D Daya Bay constraints.



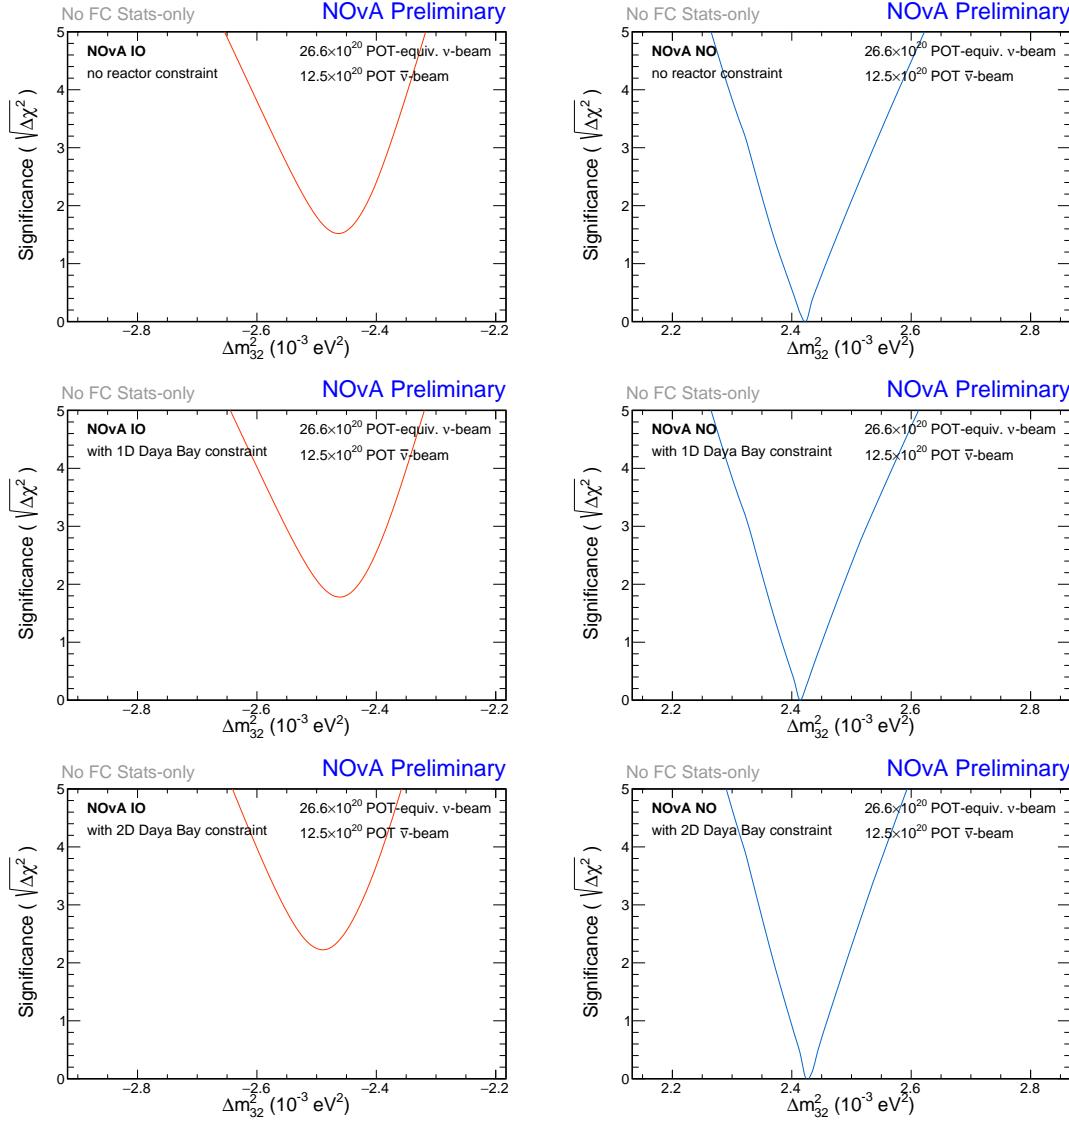
**Figure 7.14:** Feldman-Cousins corrected 68% and 90% confidence level contours for the  $\Delta m_{32}^2$  vs.  $\sin^2 \theta_{23}$  space, found from a frequentist fit to the data. The fit included the 1D Daya Bay reactor constraint and all systematic uncertainties. From [158].

Parameter	Best-fit	Normal Ordering Preference ( $\sigma$ )	
$\sin^2(\theta_{23})$	$0.546^{+0.032}_{-0.075}$	W/ 1D Daya Bay constraint	p-value 0.1731 $1.36\sigma$
$\Delta m_{32}^2 (10^{-3} \text{ eV}^2)$	$2.433^{+0.035}_{-0.036}$		
$\delta_{\text{CP}}(\pi)$	0.875	W/ 2D Daya Bay constraint	p-value 0.1158 $1.57\sigma$

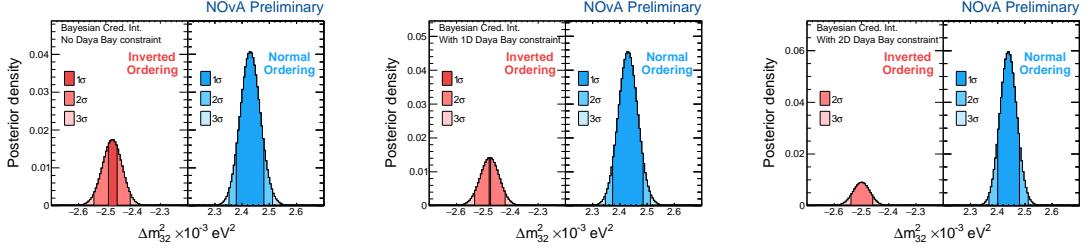
**Figure 7.15: Left:** The best-fit oscillation parameters and FC-corrected  $1\sigma$  uncertainties from the frequentist fit to the data, using the 1D Daya Bay constraint. **Right:** The significance with which we prefer the normal mass ordering when using the 1D or 2D Daya Bay constraints. From [158].

### 7.2.2 The Effect of Reactor Neutrino Constraints

Constraints from reactor neutrino experiments like Daya Bay are expected to enhance our mass ordering sensitivity, as described by Nunokawa, Parke, and Funchal [159]. Figure 7.16 shows the 1D stats-only fits to  $\Delta m_{32}^2$  using different options for the reactor constraint. Unlike the other 1D frequentist plots in this chapter, the  $\Delta\chi^2$  values in both mass ordering are calculated with respect to the global best-fit minimum in the normal mass ordering, to illustrate the difference in  $\chi^2$  between normal and inverted mass orderings. As we go from no reactor constraint, to a 1D and then 2D constraint, the difference between the minimum  $\chi^2$  in normal and inverted mass ordering increases. This can be seen as the inverted curves being displaced further upwards, showing our enhanced preference for the normal mass ordering. Figure 7.17 shows a similar result from the Bayesian framework. The plot shows the marginalized posterior probability densities for  $\Delta m_{32}^2$  using the same three options for reactor constraint (no constraint, 1D, or 2D Daya Bay constraint). Without a reactor constraint, we have a 70% preference for the normal mass ordering (i.e. 70% of the total posterior density is in the normal mass ordering). With the 1D Daya Bay constraint this increases to 77%, and with the 2D constraint it increases further to an 87% preference.

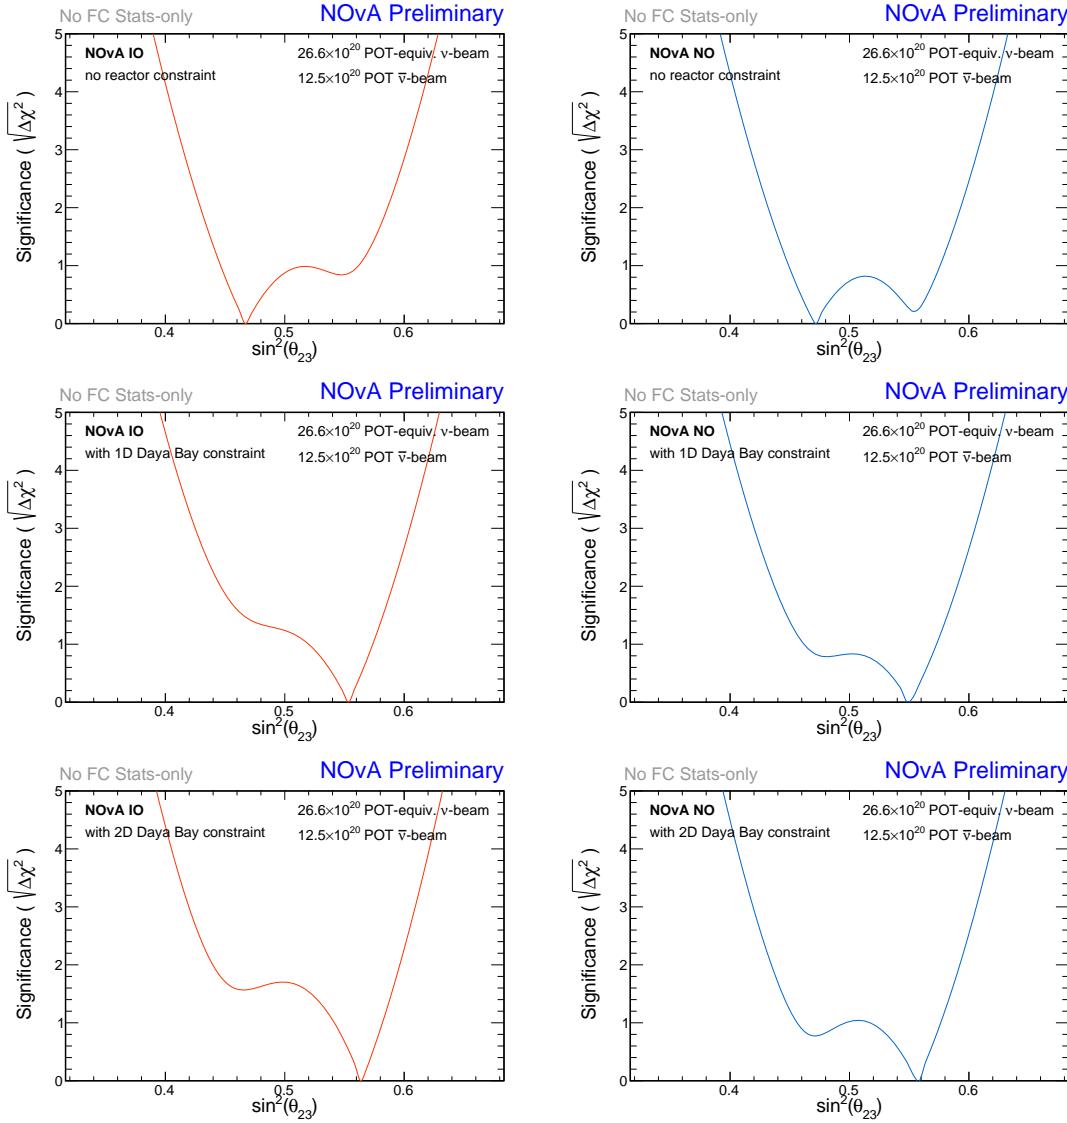


**Figure 7.16:** 1D stats-only fits of  $\Delta m_{32}^2$  using no reactor constraint (top), the 1D Daya Bay constraint (middle), or the 2D Daya Bay constraint (bottom). Here the significance is drawn relative to the global best-fit point, to highlight the increasing preference for the normal MO we get by applying the 1D and then 2D Daya Bay constraints.

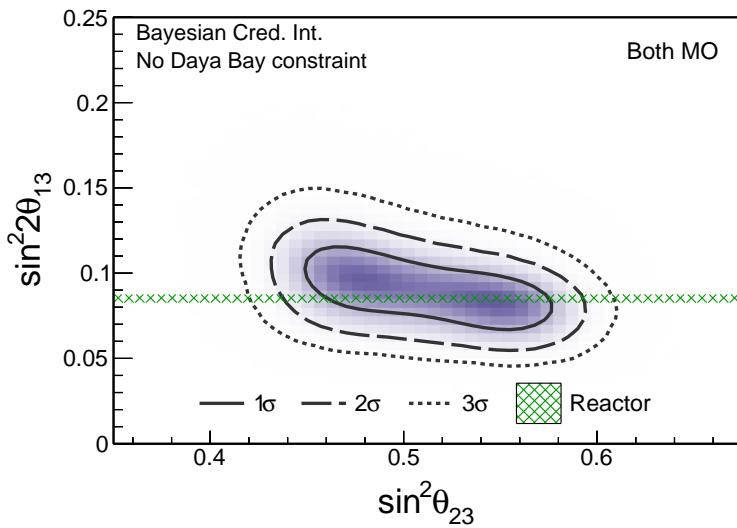


**Figure 7.17:** Marginalized posterior probability densities for  $\Delta m_{32}^2$  in each mass ordering, from the Bayesian fit with systematics included, and for three different reactor constraint options. Applying the 1D and then 2D Daya Bay constraints increases our preference for the normal mass ordering compared to the NOvA-only measurement. From [160].

The reactor constraint also serves to enhance our preference for the upper octant of  $\theta_{23}$ . Figure 7.18 shows the 1D stats-only frequentist fits for  $\sin^2 \theta_{23}$  using the three reactor constraint options. There is a slight degeneracy, with local minima in either octant. Without a reactor constraint we have a slight preference for the lower octant of  $\theta_{23}$ . Applying the 1D constraint changes the preference to the upper octant, and applying the 2D constraint enhances the preference further. It is not obvious in the 1D plots how the degeneracy in  $\sin^2 \theta_{23}$  is broken by the reactor constraint. To see how, we can look at a 2D plot of  $\sin^2 2\theta_{13}$  vs.  $\sin^2 \theta_{23}$ . Figure 7.19 shows the Bayesian posterior density and credible regions in  $\sin^2 2\theta_{13}$  vs.  $\sin^2 \theta_{23}$  from a fit without any reactor constraints (NOvA-only measurement). The  $1\sigma$  range from the 1D  $\sin^2 2\theta_{13}$  Daya Bay measurement is overlaid to illustrate the effect. There is more overlap between the Daya Bay result and the posterior density in the upper octant.



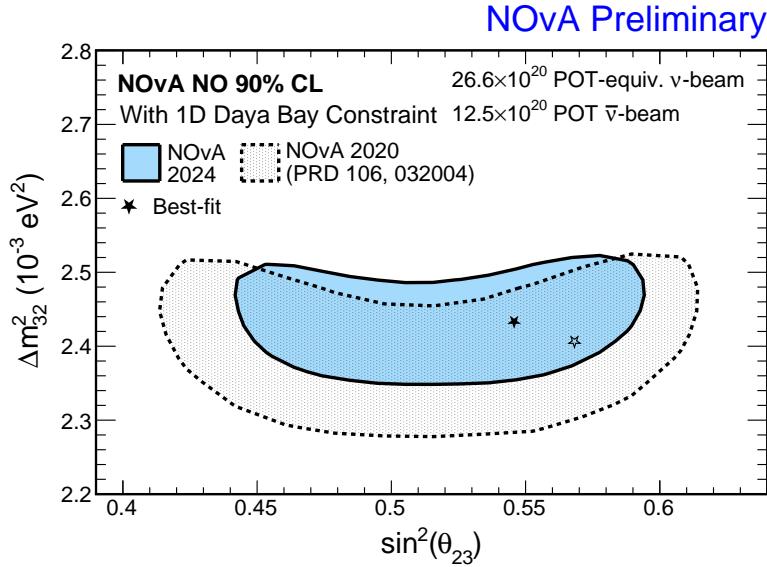
**Figure 7.18:** 1D stats-only frequentist fits of  $\sin^2 \theta_{23}$  using no reactor constraint (top), the 1D Daya Bay constraint (middle), or the 2D Daya Bay constraint (bottom). Applying the 1D constraint shifts our preference from the lower to upper octant of  $\theta_{23}$ , with the 2D constraint increasing the preference further.



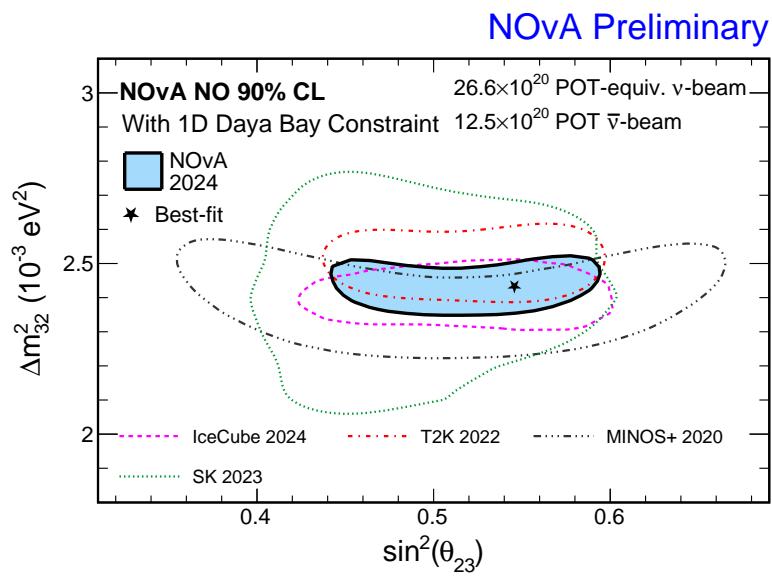
**Figure 7.19:** 2D Posterior probability density and Bayesian credible regions for  $\sin^2 \theta_{23}$  vs.  $\sin^2 2\theta_{13}$  marginalized jointly over the mass orderings. Posteriors are extracted from a fit to NOvA data without any reactor constraints, with results from the Daya Bay experiment overlaid as the green  $1\sigma$  band [139]. From [160].

### 7.2.3 NOvA Results in a Global Context

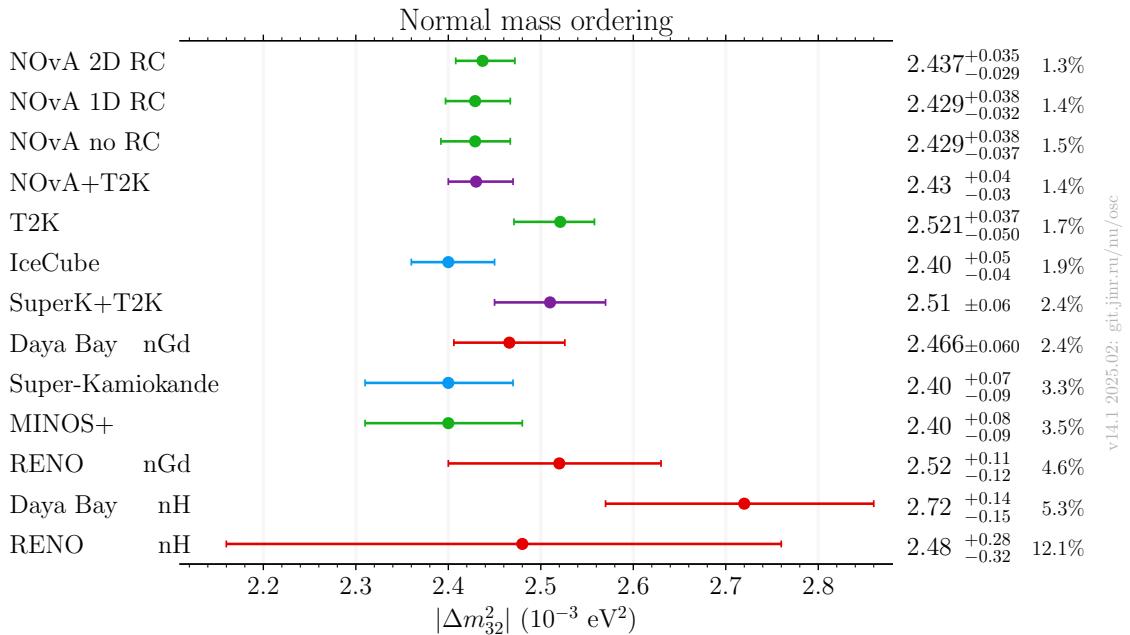
Here we compare the new results to both NOvA’s previous results, and the results of the broader neutrino oscillation physics community. Figure 7.20 shows a comparison of the FC-corrected frequentist confidence regions for  $\Delta m_{32}^2$  vs.  $\sin^2 \theta_{23}$  to those from NOvA’s 2020 analysis. The new result occupies a similar region of phase space but with tighter constraints in both directions compared to the previous result. Figure 7.21 shows a comparison of the new result to results from other neutrino oscillation experiments. We see good agreement between all results in this space. Figure 7.22 shows our central values and  $1\sigma$  ranges of  $\Delta m_{32}^2$  from the Bayesian analysis, compared to those from other experiments. With this new result, NOvA currently has the most precise single-experiment measurement of  $\Delta m_{32}^2$ , as shown in Fig. 7.22.



**Figure 7.20:** Comparison of the  $\Delta m_{32}^2$  vs.  $\sin^2 \theta_{23}$  FC-corrected 90% confidence level contour from this analysis to the one from NOvA’s 2020 analysis [113]. From [158].

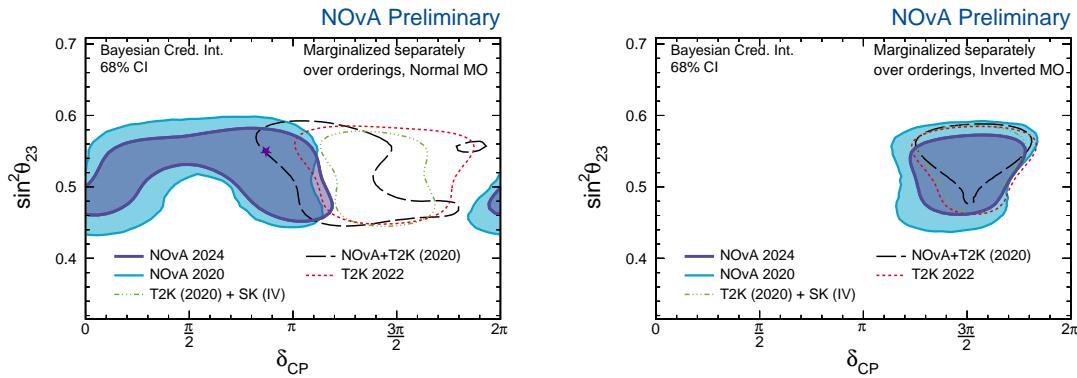


**Figure 7.21:** Comparisons of the 90% confidence level contours for  $\Delta m_{32}^2$  in the normal MO between our new result and recent results from other experiments [161–164]. From [158].



**Figure 7.22:** Comparison of central values and  $1\sigma$  intervals for  $\Delta m_{32}^2$  in the normal MO from various accelerator (green), reactor (red), and atmospheric (blue) experiments, and two joint fits (purple). The exact values are shown on the right along with the % uncertainty. The first three NOvA results are found from the Bayesian analysis with different reactor constraints. At 1.5% uncertainty, the NOvA no RC result is currently the most precise single-experiment measurement of  $\Delta m_{32}^2$ . Sources for the remaining results, beginning with the fourth line are as follows: [139, 161–168]. From [169].

In the  $\sin^2 \theta_{23}$  vs.  $\delta_{\text{CP}}$  space NOvA results are again consistent with our 2020 results, but with improved sensitivity. However, there is a small tension with the results from T2K in the normal MO that has persisted since the 2020 analysis. Figure 7.23 shows the Bayesian 68% credible regions in  $\sin^2 \theta_{23}$  vs.  $\delta_{\text{CP}}$  for recent NOvA and T2K results. In the normal MO NOvA prefers a region of  $\delta_{\text{CP}}$  that leads to low asymmetry in  $\nu_e$  ( $\bar{\nu}_e$ ) appearance rates, while T2K prefers a large asymmetry. In the inverted mass ordering there is good agreement between experiments. Currently the difference between the experiments does not occur at a high degree of statistical significance, with more overlap at 2 and  $3\sigma$  levels. If the tension is to be resolved, additional datataking, particularly in antineutrino mode, will be necessary.



**Figure 7.23:** New NOvA results for  $\sin^2 \theta_{23}$  vs.  $\delta_{\text{CP}}$  in the Normal (left) and Inverted (right) MOs, compared with previous NOvA results [113, 138] and results from other experiments [161, 164, 170], including the 2024 joint NOvA-T2K analysis [171, 172]. All of the credible intervals are extracted from Bayesian analyses. The star indicates NOvA’s highest probability density point from this analysis. From [169].

# Chapter 8

## Conclusions

This thesis presents an analysis of  $26.61 \times 10^{20}$  POT of neutrino-mode and  $12.5 \times 10^{20}$  POT of antineutrino-mode beam data. On top of doubling our FHC dataset, this analysis included improvements to the light-model tune, improved systematic uncertainties, a new sample of  $\nu_e$  events, and an additional focus on reactor neutrino constraints including our first implementation of a constraint on  $\Delta m_{32}^2$ . In the muon neutrino disappearance channel we get sensitivity to  $\Delta m_{32}^2$  and  $\sin^2 \theta_{23}$  primarily through measurement of the minimum of the oscillation dip region of the energy spectra. The frequentist fit to the data, using the 1D Daya Bay constraint on  $\sin^2 2\theta_{13}$ , and with Feldman-Cousins corrections applied, yielded best-fit values of  $\Delta m_{32}^2 = 2.433_{-0.036}^{+0.035}$  ( $10^{-3}$  eV $^2$ ), and  $\sin^2 \theta_{23} = 0.546_{-0.075}^{+0.032}$ . Without a reactor constraint, the Bayesian analysis finds  $\Delta m_{32}^2 = 2.429_{-0.037}^{+0.038}$  ( $10^{-3}$  eV $^2$ ) which at 1.5% uncertainty is currently the most accurate single-experiment measurement of  $\Delta m_{32}^2$ . In the electron neutrino appearance channel we get sensitivity to the mass ordering, value of  $\delta_{\text{CP}}$ , and octant of  $\theta_{23}$  through measurement of asymmetry in the rates of  $\nu_e$  ( $\bar{\nu}_e$ ) appearance. Our best-fit value falls in the normal mass ordering and upper-octant of  $\theta_{23}$ , with  $\delta_{\text{CP}} = 0.875\pi$ . With the 1D Daya Bay constraint applied we have a  $1.36\sigma$  preference for the normal mass ordering, which increases to  $1.57\sigma$  when applying the 2D Daya Bay constraint. Our data prefers combinations of oscillation parameters that lead to little asymmetry in the  $\nu_e$  ( $\bar{\nu}_e$ ) appearance channels. We disfavor regions of high asymmetry,

excluding values of  $\delta_{\text{CP}}$  from  $0\text{-}\pi$ , and  $\delta_{\text{CP}} = 2\pi$ , in the inverted mass ordering at  $>3\sigma$ . This puts us in a low-significance tension with T2K [173], who prefer oscillation parameters that lead to significant asymmetry in the appearance channel. Along with this mild tension, there are many open questions that warrant further analysis. The true mass ordering, octant of  $\theta_{23}$ , and value of  $\delta_{\text{CP}}$  are still not known at a high significance level.

NOvA has an additional  $2.8 \times 10^{20}$  POT of antineutrino data waiting to be analyzed, and hopes to take an additional  $1.5\text{-}2.5 \times 10^{20}$  POT of antineutrino data before concluding operations at the beginning of 2027. This is to enable the beginning a long shutdown period in preparation for DUNE, which will make its first measurements early next decade. In this interim period, NOvA will release it's final 3-flavor oscillation results using our full dataset. This will be among the last new long-baseline results for the community until the next-generation of experiments such as DUNE come online. Further analysis improvements are being explored to ensure we get the maximum sensitivity possible out of this measurement. NOvA has data from a test-beam run that are being analyzed to improve our large energy scale uncertainties. Updated simulation models and reconstruction techniques have been incorporated into the production of our final dataset. Some of the studies to reclaim neutrino events that were conducted for this analysis can then be re-examined with more data and re-tuned selections. The improvements to Michel electron reconstruction will be fully incorporated as well, and explored further for potential sensitivity improvements.

Along with our independent results, collaboration with other experiments will be critical. NOvA and T2K have performed a joint 3-flavor analysis that is being prepared for publication [171, 172]. We continue to collaborate and plan further joint analyses that can exploit the unique qualities of each experiment to break degeneracies, and hopefully resolve the tension between them. The next-generation reactor neutrino experiment JUNO is expected to begin taking data this year, and will likely be the first of the next-generation experiments to publish results. If NOvA's final 3-flavor results can push the uncertainty on  $\Delta m_{32}^2$  to the 1% level, then we could combine our results with JUNO's high-precision measurement of  $\Delta m_{32}^2$  to make an even better measurement of the mass ordering [159].

# References

- [1] J. Chadwick. The intensity distribution in the magnetic spectrum of  $\beta$  particles from radium (B + C). *Verh. Phys. Gesell.*, 16:383–391, 1914.
- [2] E. Fermi. Versuch einer theorie der  $\beta$ -strahlen. i. *I. Z. Physik*, 88, March 1934.
- [3] B. Pontecorvo. Inverse beta process. *Camb. Monogr. Part. Phys. Nucl. Phys. Cosmol.*, 1:25–31, 1991.
- [4] C. Cowan, F. Reines, F. Harrison, H. Kruse, and A. McGuire. Detection of the free neutrino: a confirmation. *Science*, 124, 1956.
- [5] G. Danby et al. Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos. *Phys. Rev. Lett.*, 9, 1962.
- [6] Martin L. Perl et al. Evidence for Anomalous Lepton Production in e+ - e- Annihilation. *Phys. Rev. Lett.*, 35:1489–1492, 1975.
- [7] T. Patzak (DONUT). First direct observation of the tau-neutrino. *Europhys. News*, 32:56–57, 2001.
- [8] B. Pontecorvo. Mesonium and Antimesonium. *Sov. Phys. JETP*, 6:429–431, 1958.
- [9] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. Remarks on the unified model of elementary particles. *Prog. Theor. Phys.*, 28:870–880, 1962.

[10] B. Pontecorvo. Neutrino Experiments and the Problem of Conservation of Leptonic Charge. *Zh. Eksp. Teor. Fiz.*, 53:1717–1725, 1967.

[11] Raymond Davis, Don S. Harmer, and Kenneth C. Hoffman. Search for neutrinos from the sun. *Phys. Rev. Lett.*, 20:1205–1209, May 1968.

[12] Y. Fukuda et al. (Super-Kamiokande). Evidence for oscillation of atmospheric neutrinos. *Phys. Rev. Lett.*, 81:1562–1567, 1998.

[13] Q. Ahmad et al. (SNO Collaboration). Direct evidence for neutrino flavor transformation from neutral-current interactions in the sudbury neutrino observatory. *Phys. Rev. Lett.*, 89:011301, Jun 2002.

[14] Wikimedia Commons. Standard model of elementary particlese. [Web image](#), 2019. Accessed: 2025-7-27.

[15] Fernanda Psihas. *Measurement of Long Baseline Neutrino Oscillations and Improvements from Deep Learning*. PhD thesis, Indiana U., 2018.

[16] Alexander Craig Booth. *Electron neutrino appearance at the NOvA experiment*. PhD thesis, Univ. Of Sussex, 10 2021.

[17] Pierre Lasorak. *A search for neutrino-induced single photons and measurement of oscillation analysis systematic errors with electron and anti-electron neutrino selections, using the off-axis near detector of the Tokai to Kamioka experiment*. PhD thesis, Queen Mary, U. of London, 2018.

[18] Miranda Rabelhofer. *Study of neutron detector response and related systematic uncertainties in the NOvA oscillation analysis*. PhD thesis, Iowa State U., 2023.

[19] J. A. Formaggio and G. P. Zeller. From ev to eev: Neutrino cross sections across energy scales. *Rev. Mod. Phys.*, 84:1307–1341, Sep 2012.

[20] P. Hernandez. Neutrino Physics. In *8th CERN–Latin-American School of High-Energy Physics*, pages 85–142, 2016. 10.5170/CERN-2016-005.85.

[21] L. Wolfenstein. Neutrino Oscillations in Matter. *Phys. Rev. D*, 17:2369–2374, 1978.

[22] S. P. Mikheyev and A. Yu. Smirnov. Resonance Amplification of Oscillations in Matter and Spectroscopy of Solar Neutrinos. *Sov. J. Nucl. Phys.*, 42:913–917, 1985.

[23] S. Navas et al. (Particle Data Group). Review of particle physics. *Phys. Rev. D*, 110(3):030001, 2024.

[24] Peter B. Denton and Stephen J. Parke. Smallness of matter effects in long-baseline muon neutrino disappearance. *Phys. Rev. D*, 109:053002, Mar 2024.

[25] Hiroshi Nunokawa, Stephen J. Parke, and Jose W. F. Valle. CP Violation and Neutrino Oscillations. *Prog. Part. Nucl. Phys.*, 60:338–402, 2008.

[26] Mattias Blennow and Alexei Yu. Smirnov. Neutrino propagation in matter. *Adv. High Energy Phys.*, 2013:972485, 2013.

[27] I. Singh. New Perspectives - Three-Flavor Neutrino Oscillations at NOvA. NOvA Internal Documents [docid=63057](#), 2024.

[28] Nitish Nayak. Pedagogical Bi-Probability Plots – Blessing Package. NOvA Internal Documents [docid=38454](#), 2019.

[29] Jeremy Wolcott. New neutrino oscillation results from nova with 10 years of data. Zenodo. <https://doi.org/10.5281/zenodo.12704805>, June 2024.

[30] T. Lackey. US Map with NOvA Baseline . NOvA Internal Documents [docid=66945](#), 2025.

[31] Katsuya Yonehara, Sudeshna Ganguly, Don Athula Wickremasinghe, Pavel Snopok, and Yiding Yu. Exploring the Focusing Mechanism of the NuMI Horn Magnets. 5 2023.

[32] Accelerator Division Operations Department. *Concepts Rookie Book*. Tech. Rep. (Fermilab, 2020).

[33] Fermilab Creative Services. 'fcs still photos'. <https://mod.fnal.gov>. Accessed Jan 1st, 2025.

[34] Katsuya Yonehara (NuMI). Megawatt upgrade of NuMI target system. *PoS, NuFact2021:107*, 2022.

[35] P. Adamson et al. The NuMI neutrino beam. *Nucl. Instrum. Methods Phys Res., Sect. A*, 806:279–306, 2016.

[36] J. Hylen. NuMI Status . NOvA Internal Documents [docid=29534](#), 2018.

[37] S. Kopp et al. Secondary beam monitors for the numi facility at fnal. *Nucl. Instrum. Meth. A*, 568(2):503–519, December 2006.

[38] D. S. Ayres et al. (NOvA). *The NOvA Technical Design Report*. October 2007.

[39] M. D. Tutto. Blessing Package - Beam Simulation Plots. NOvA Internal Documents [docid=13524](#), 2015.

[40] Erika Catano-Mur. *Constraints on neutrino oscillation parameters with the NOvA experiment*. PhD thesis, Iowa State University, 2018.

[41] M. Rabelhofer. Ben+Michel Decomposition Blessing Package 2020. NOvA Internal Documents [docid=45865](#), 2020.

[42] R.L. Talaga et al. Pvc extrusion development and production for the nova neutrino experiment. *Nucl. Instrum. Meth. A*, 861:77–89, 2017.

[43] The NOvA APD is a custom variant of the Hamamatsu S8550 <https://www.hamamatsu.com/us/en/product/optical-sensors/apd/si-apd-array/S8550-02.html>.

[44] A. Salles. Press Release: Fermilab’s 500-mile neutrino experiment up and running. <https://news.fnal.gov/2014/10/fermilabs-500-mile-neutrino-experiment-up-and-running/>, 2014.

[45] A. Norman (NOvA). The NOvA Data Acquisition System. *J. Phys. Conf. Ser.*, 396:012035, 2012.

[46] J. Zalesak. The NOvA Far Detector Data Acquisition System. NOvA Internal Documents [docid=10032](#), 2013.

[47] M. A. Acero et al. (NOvA). Supernova neutrino detection in NOvA. *JCAP*, 10:014, 2020.

[48] E. Niner. Timing Calibration Technical Note. NOvA Internal Documents [docid=12570](#), 2015.

[49] D. Bhattacharai. APS 2022 Data Quality Talk: Blessing Package. NOvA Internal Documents [docid=54441](#), 2022.

[50] A. Cooleybeck. Blessing package for FD All POT plots. NOvA Internal Documents [docid=62786](#), 2024.

[51] R. Hatcher. Simulation Tools in Neutrino Experiments. NOvA Internal Documents [docid=14441](#), 2013.

[52] Júlia Tena-Vidal et al. Neutrino-Nucleon Cross-Section Model Tuning in GENIE v3. FERMILAB-PUB-20-531-SCD-T, 4 2021.

[53] Geant4 Collaboration. Geant4 10.4 release notes. [geant4-data.web.cern.ch](http://geant4-data.web.cern.ch/), <https://geant4-data.web.cern.ch/ReleaseNotes/ReleaseNotes4.10.4.html>, 2017.

[54] L. Aliaga. 2017-2018 Beam Plots. NOvA Internal Documents [docid=20843](#), 2018.

[55] L. Aliaga et al. (MINERvA). Neutrino Flux Predictions for the NuMI Beam. *Phys. Rev. D*, 94(9):092005, 2016. [Addendum: Phys.Rev.D 95, 039903 (2017)].

[56] C. Alt et al. (NA49). Inclusive production of charged pions in p+C collisions at 158-GeV/c beam momentum. *Eur. Phys. J. C*, 49:897–917, 2007.

[57] J. Wolcott. Recent Cross Section Work From NOvA. In *18th International Workshop on Neutrino Factories and Future Neutrino Facilities Search*, 11 2016.

[58] M. Martinez Casales. *Constraining neutrino interaction uncertainties for oscillation measurements in the NOvA experiment using Near Detector data*. PhD thesis, Iowa State U. (main), Iowa State U., 2023.

[59] M. Martinez Casales. Ana 2024 (prod5.1) MEC tuning blessing package. NOvA Internal Documents [docid=62915](#), 2024.

[60] J. Nieves, J. E. Amaro, and M. Valverde. Inclusive quasi-elastic neutrino reactions. *Phys. Rev. C*, 70:055503, 2004. [Erratum: Phys.Rev.C 72, 019902 (2005)].

[61] J. Nieves et al. Inclusive charged-current neutrino-nucleus reactions. *Phys. Rev. C*, 83:045501, Apr 2011.

[62] R. Gran, J. Nieves, F. Sanchez, and M. J. Vicente Vacas. Neutrino-nucleus quasi-elastic and 2p2h interactions up to 10 GeV. *Phys. Rev. D*, 88(11):113007, 2013.

[63] Ch. Berger and L. M. Sehgal. Lepton mass effects in single pion production by neutrinos. *Phys. Rev. D*, 76:113004, Dec 2007.

[64] Aaron S. Meyer, Minerba Betancourt, Richard Gran, and Richard J. Hill. Deuterium target data for precision neutrino-nucleus cross sections. *Phys. Rev. D*, 93:113015, Jun 2016.

[65] A. Bodek and U. K. Yang. Modeling deep inelastic cross-sections in the few GeV region. *Nucl. Phys. B Proc. Suppl.*, 112:70–76, 2002.

[66] Costas Andreopoulos, Christopher Barry, Steve Dytman, Hugh Gallagher, Tomasz Golan, Robert Hatcher, Gabriel Perdue, and Julia Yarba. The GENIE Neutrino Monte Carlo Generator: Physics and User Manual. arXiv 1510.05494, June 2015.

[67] B. W. Allardice et al. Pion reaction cross-sections and nuclear sizes. *Nucl. Phys. A*, 209:1–51, 1973.

[68] A. Saunders, S. Hoeibraten, J. J. Kraushaar, B. J. Kriss, R. J. Peterson, R. A. Ristinen, J. T. Brack, G. Hofman, E. F. Gibson, and C. L. Morris. Reaction and total cross-sections for low-energy  $\pi^+$  and  $\pi^-$  on isospin zero nuclei. *Phys. Rev. C*, 53:1745–1752, 1996.

[69] O. Meirav, E. Friedman, R. R. Johnson, R. Olszewski, and P. Weber. Low-energy Pion - Nucleus Potentials From Differential and Integral Data. *Phys. Rev. C*, 40:843–849, 1989.

[70] S. M. Levenson et al. Inclusive pion scattering in the delta (1232) region. *Phys. Rev. C*, 28:326–332, 1983.

[71] D. Ashery, I. Navon, G. Azuelos, H. K. Walter, H. J. Pfeiffer, and F. W. Schleputz. True Absorption and Scattering of Pions on Nuclei. *Phys. Rev. C*, 23:2173–2185, 1981.

[72] D. Ashery et al. Inclusive pion single charge exchange reactions. *Phys. Rev. C*, 30:946–951, 1984.

[73] E. S. Pinzon Guerra et al. (DUET). Measurement of  $\sigma_{\text{ABS}}$  and  $\sigma_{\text{CX}}$  of  $\pi^+$  on carbon by the Dual Use Experiment at TRIUMF (DUET). *Phys. Rev. C*, 95(4):045203, 2017.

[74] O. Samoylov. Birks fit validation for results from Dubna scintillator stand measurement. NOvA Internal Documents [docid=181120](#), 2018.

[75] A. Aurisano, C. Backhouse, R. Hatcher, N. Mayer, J. Musser, R. Patterson, R. Schroeter, and A. Sousa (NOvA). The NOvA simulation chain. *J. Phys. Conf. Ser.*, 664(7):072002, 2015.

[76] P. Balaban, P.E. Fleischer, and H. Zucker. The probability distribution of gains in avalanche photodiodes. *IEEE Transactions on Electron Devices*, 23(10):1189–1190, 1976.

[77] Chris Hagmann, David Lange, and Douglas Wright. Cosmic-ray shower generator (cry) for monte carlo transport codes. In *2007 IEEE Nuclear Science Symposium Conference Record*, volume 2, pages 1143–1146, 2007.

[78] Tyler Alion. Third Analysis: Absolute Calibration Technote. NOvA Internal Documents [docid=23372](#), 2017.

[79] R. L. Workman et al. (Particle Data Group). Review of Particle Physics. *PTEP*, 2022:083C01, 2022.

[80] FD Prod5.1 Calibration Validation. L. Koerner and J. Lesmeister. NOvA Internal Documents [docid=54054](#), 2022.

[81] L. Vinton. Calorimetric Energy Scale Calibration of the NOvA Detectors. NOvA Internal Documents [docid=13579](#), 2015.

[82] C. Backhouse, A. Radovic, P. Singh, and M. Campbell. The Attenuation and Threshold Calibration of the NOvA detectors. NOvA Internal Documents [docid=13579](#), 2017.

[83] A. Lister and J. Trokan-Tenorio. Calibration of Periods 12-14 Data in the NOvA Near and Far Detectors . NOvA Internal Documents [docid=60838](#), 2024.

[84] C. Green, J. Kowalkowski, M. Paterno, M. Fischler, L. Garren, and Q. Lu. The Art Framework. *J. Phys. Conf. Ser.*, 396:022020, 2012.

[85] F. Psihas. Event Displays for Nue Selected Events. NOvA Internal Documents [docid=15647](#), 2016.

[86] Alex Rodriguez and Alessandro Laio. Clustering by fast search and find of density peaks. *Science*, 344(6191):1492–1496, 2014.

[87] D. Pershey, J. Huang, and M. Judah. TDSlicer Technote. NOvA Internal Documents [docid=27689](#), 2019.

[88] R. C. Prim. Shortest connection networks and some generalizations. *The Bell System Technical Journal*, 36(6):1389–1401, 1957.

[89] M. Baird, J. Bian, M. Messier, E. Niner, D. Rocco, and K. Sachdev. Event reconstruction techniques in nova. *Journal of Physics: Conference Series*, 664(7):072035, dec 2015.

[90] Leandro A.F. Fernandes and Manuel M. Oliveira. Real-time line detection through an improved hough transform voting scheme. *Pattern Recognition*, 41(1):299–314, 2008.

[91] Mattias Ohlsson. Extensions and explorations of the elastic arms algorithm. *Computer Physics Communications*, 77(1):19–32, 1993.

[92] R. Raddatz. KalmanTrack Technical Note. NOvA Internal Documents [docid=13545](#), 2015.

[93] R. E. Kalman. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 82(1):35–45, 03 1960.

[94] B. Rebel. A Window Tracking Algorithm for Cosmic Ray Muons. NOvA Internal Documents [docid=15997](#), 2016.

[95] D. Pershey. MEFinder Technote. NOvA Internal Documents [docid=14789](#), 2016.

[96] Martin Ester, Hans-Peter Kriegel, Jörg Sander, and Xiaowei Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, KDD'96, page 226–231. AAAI Press, 1996.

[97] A. Aurisano, A. Radovic, D. Rocco, A. Himmel, M.D. Messier, E. Niner, G. Pawloski, F. Psihas, A. Sousa, and P. Vahle. A convolutional neural network neutrino event classifier. *Journal of Instrumentation*, 11(09):P09001, sep 2016.

[98] M. Groh and K. Warburton. CVN Training for the 2020 Analysis. NOvA Internal Documents [docid=42897](#), 2020.

[99] F. Psihas. Assorted CVN Plots for Blessing. NOvA Internal Documents [docid=15639](#), 2017.

[100] J. Porter. ReMId 2019 retraining. NOvA Internal Documents [docid=42277](#), 2019.

[101] P. Gandrajula and L. Kolupaeva. Cosmic rejection BDT for e appearance analysis. NOvA Internal Documents [docid=42473](#), 2019.

[102] K. Warburton. Muon Neutrino Cosmic Rejection Technote For The 2020 Analysis. NOvA Internal Documents [docid=42358](#), 2020.

[103] A. Back, C. Sullivan, and M. Messier. A new low-energy classifier to enhance nue selection. NOvA Internal Documents [docid=59554](#), 2023.

[104] W. Wu and A. Kalitkina. Prod5.1 Nue Energy Estimator. NOvA Internal Documents [docid=61091](#), 2024.

[105] E. Catano-Mur, R. Nichol, and Z. Vallari. Three-Flavour Executive Summary 2024 . NOvA Internal Documents [docid=61655](#), 2024.

[106] B. Jargowsky and R. Nichol. Prod5.1 Numu Energy Estimators Technote. NOvA Internal Documents [docid=61064](#), 2024.

[107] B. Jargowsky. Blessed Package For Prod5.1 Energy Estimation for 3F Ana2024. NOvA Internal Documents [docid=61952](#), 2024.

[108] J. Coelho, B. Chowdhury, and R. Murphy. Good Data Selection. NOvA Internal Documents [docid=13546](#), 2017.

[109] S. Lein. DCM Edge Metric. NOvA Internal Documents [docid=13527](#), 2015.

[110] Ashley Back et al. Data Quality 2020. NOvA Internal Documents [docid=44226](#), 2020.

[111] Michael Baird, Louise Suter, and Jeremy Wolcott. Summary of the 2020 FHC+RHC 3-Flavor oscillation analysis. NOvA Internal Documents [docid=44422](#), 2020.

[112] E. Catano-Mur et al. 3-Flavour event selection for Ana2024 analysis. NOvA Internal Documents [docid=61640](#), 2024.

[113] M. A. Acero et al. (NOvA). Improved measurement of neutrino oscillation parameters by the NOvA experiment. *Phys. Rev. D*, 106(3):032004, 2022.

[114] M. Baird and L. Vinton. Extrapolation Technote for the Numu Third Analysis. NOvA Internal Documents [docid=23390](#), 2017.

[115] W. Shorrock. Numu Binning Studies. NOvA Internal Documents [docid=60579](#), 2024.

[116] A. Mislevic. Near-to-Far Extrapolation in  $p_T$  for the NOA 2020 3-Flavor Analysis. NOvA Internal Documents [docid=44401](#), 2020.

[117] M. Rabelhofer. 2020 Decomposition: Michel, BEN, and Proportional. NOvA Internal Documents [docid=43712](#), 2020.

[118] E. Catano-Mur. BEN Decomposition Tech Note. NOvA Internal Documents [docid=15392](#), 2016.

[119] D. Pershey. MichelDecomp Technote for the 2017  $\nu_e$  Analysis. NOvA Internal Documents [docid=22523](#), 2017.

[120] M. Rabelhofer. LowE  $\nu_e$  Extrapolation Technote. NOvA Internal Documents [docid=62120](#), 2024.

[121] J. Lozier. ModularExtrap Technical Note. NOvA Internal Documents [docid= 12563](#), 2014.

[122] T. Bezerra, E. Catano-Mur, and R. Nichol. Cosmic predictions for 3F 2024 analysis. NOvA Internal Documents [docid=62105](#), 2024.

[123] N. Nayak. Hadron Production Systematics for the NOvA Oscillation Analysis. NOvA Internal Documents [docid=22532](#), 2017.

[124] N. Nayak. Flux Systematics for the 2018 NOvA Oscillation Analyses. NOvA Internal Documents [docid=27884](#), 2018.

[125] Kirk Bays et al. NOvA Cross-Section Modeling Internal Technical Note For Production 5.1. NOvA Internal Documents [docid=61559](#), 2024.

[126] B. Jargowsky. 2024 Genie PCA Technote. NOvA Internal Documents [docid=62099](#), 2024.

[127] Erika Catano-Mur, V Hewes, and Lisa Koerner. Prod 5.1 Detector Systematics Executive Summary. NOvA Internal Documents [docid=53225](#), 2022.

[128] M. Strait. Muon energy scale systematic. NOvA Internal Documents [docid=20816](#), 2022.

[129] J. Calcutt, C. Thorpe, K. Mahn, and L. Fields. Geant4reweight: a framework for evaluating and propagating hadronic interaction uncertainties in geant4. *Journal of Instrumentation*, 16(08):P08042, aug 2021.

[130] B. Jargowsky. 2024 Geant4Reweight PCA Technote. NOvA Internal Documents [docid=62099](#), 2024.

[131] B Roeder. Development and validation of neutron detection simulations for eurisol. *EURISOL Design Study, Report No. 10-25-2008-006-In-beamvalidations.pdf*, pp 31-44, 2008.

[132] Z. Kohley et al. Modeling interactions of intermediate-energy neutrons in a plastic scintillator array with Geant4. *Nucl. Instrum. Meth. A*, 682:59–65, 2012.

[133] M. Rabelhofer. File-based MENATE Neutron Systematic. NOvA Internal Documents [docid=59999](#), 2024.

[134] A. Yankelevich. Ana2024: 3F Normalization Systematics. NOvA Internal Documents [docid=62214](#), 2024.

[135] J. Trokan-Tenorio. Michel Electron Tagging Systematic Overview For Ana2024. NOvA Internal Documents [docid=62217](#), 2024.

[136] M. A. Acero et al. (NOvA). New constraints on oscillation parameters from  $\nu_e$  appearance and  $\nu_\mu$  disappearance in the NOvA experiment. *Phys. Rev. D*, 98:032012, 2018.

[137] M. A. Acero et al. (NOvA). First Measurement of Neutrino Oscillation Parameters using Neutrinos and Antineutrinos by NOvA. *Phys. Rev. Lett.*, 123(15):151803, 2019.

[138] M. A. Acero et al. (NOvA). Expanding neutrino oscillation parameter measurements in NOvA using a Bayesian approach. *Phys. Rev. D*, 110(1):012005, 2024.

[139] F. P. An et al. (Daya Bay). Precision Measurement of Reactor Antineutrino Oscillation at Kilometer-Scale Baselines by Daya Bay. *Phys. Rev. Lett.*, 130(16):161802, 2023.

[140] Gary J. Feldman and Robert D. Cousins. A Unified approach to the classical statistical analysis of small signals. *Phys. Rev. D*, 57:3873–3889, 1998.

[141] M. A. Acero et al. (NOvA). A Monte Carlo method for constructing more accurate confidence intervals in the presence of nuisance parameters that violate Wilks' theorem assumptions. *Submitted to Phys. Rev. D*, July 2022. arXiv 2207.14353.

[142] A. Sousa, N. Buchanan, S. Calvez, P. Ding, D. Doyle, A. Himmel, B. Holzman, J. Kowalkowski, A. Norman, and T. Peterka. Implementation of feldman-cousins corrections and oscillation calculations in the hpc environment for the nova experiment. *EPJ Web Conf.*, 214:05012, 2019.

[143] A. Yankelevich et al. Ana2024: Bayesian Sensitivity. NOvA Internal Documents [docid=62090](#), 2024.

[144] A. Back, P. Lasorak, A. Sztuc, and J. Wolcott. Introduction to Bayesian MCMC fitting. NOvA Internal Documents [docid=53021](#), 2021.

[145] Benjamin Jargowsky. *A Measurement of  $\nu_e$  Appearance and  $\nu_\mu$  Disappearance Using 10 Years of Data from the NOvA Experiment*. PhD thesis, UC Irvine, 2024.

[146] Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6):1087–1092, 06 1953.

[147] Jaewook Lee, Woosuk Sung, and Joo-Ho Choi. Metamodel for efficient estimation of capacity-fade uncertainty in li-ion batteries for electric vehicles. *Energies*, 8(6):5538–5554, 2015.

[148] D. Adey et al. (The Daya Bay Collaboration). Measurement of the electron antineutrino oscillation with 1958 days of operation at daya bay. *Phys. Rev. Lett.*, 121, Dec 2018.

[149] Feng Peng An et al. (Daya Bay). Measurement of electron antineutrino oscillation based on 1230 days of operation of the Daya Bay experiment. *Phys. Rev. D*, 95(7):072006, 2017.

[150] M. Tanabashi et al. (Particle Data Group). Review of Particle Physics. *Phys. Rev. D*, 98(3):030001, 2018 and 2019 update.

[151] L. Kolupaeva. Future 3 Flavour Sensitivities - 2020 analysis. NOvA Internal Documents [docid=46343](#), 2020.

[152] Veera Mikola. *Improving the NOvA 3-Flavour Neutrino Oscillation Analysis*. PhD thesis, University College London, 2023.

[153] J. Trokan-Tenorio and C. Sullivan. Ana2024 Low-Energy Nue Plots Blessing Package. NOvA Internal Documents [docid=62876](#), 2024.

[154] D. Pershey. APD Dead Time Follow-Up. NOvA Internal Documents [docid=12295](#), 2014.

[155] J. Trokan-Tenorio. Remaking MID Templates and improving MEFinder. NOvA Internal Documents [docid=56642](#), 2022.

[156] J. Trokan-Tenorio. Miniprod6.1 Michel electrons - applying Cuts/Weights in CAFAna. NOvA Internal Documents [docid=64237](#), 2024.

[157] B. Choudhary, I. Singh, and L. Suter. Ana2024: Uncertainties on Frequentist Best-fit Point. NOvA Internal Documents [docid=62867](#), 2024.

[158] A. Dye and L. Ricardo Prais. Ana2024 Neutrino2024 Frequentist Blessing Package. NOvA Internal Documents [docid=62888](#), 2024.

[159] Hiroshi Nunokawa, Stephen Parke, and Renata Zukanovich Funchal. Another possible way to determine the neutrino mass hierarchy. *Phys. Rev. D*, 72:013009, Jul 2005.

[160] B. Jargowsky, L. Kolupaeva, and A. Yankelevich. Blessing Package For Ana2024 3F Bayesian Results. NOvA Internal Documents [docid=62918](#), 2024.

[161] K. Abe et al. First joint oscillation analysis of Super-Kamiokande atmospheric and T2K accelerator neutrino data. arXiv 2405.12488, May 2024.

[162] R. Abbasi et al. (IceCube). Measurement of Atmospheric Neutrino Oscillation Parameters Using Convolutional Neural Networks with 9.3 Years of Data in IceCube DeepCore. *Phys. Rev. Lett.*, 134(9):091801, 2025.

[163] P. Adamson et al. (MINOS+). Precision Constraints for Three-Flavor Neutrino Oscillations from the Full MINOS+ and MINOS Dataset. *Phys. Rev. Lett.*, 125(13):131802, 2020.

[164] T. Wester et al. (Super-Kamiokande). Atmospheric neutrino oscillation analysis with neutron tagging and an expanded fiducial volume in Super-Kamiokande I–V. *Phys. Rev. D*, 109(7):072014, 2024.

[165] Denis Carabadjac. T2K Neutrino Oscillation Analysis. Presented at ICHEP 2024, July 2024.

[166] S. Jeon et al. (RENO). Measurement of reactor antineutrino oscillation amplitude and frequency using 3800 days of complete data sample of the RENO experiment. December 2024. arXiv 2412.18711.

[167] F. P. An et al. (Daya Bay). Measurement of Electron Antineutrino Oscillation Amplitude and Frequency via Neutron Capture on Hydrogen at Daya Bay. *Phys. Rev. Lett.*, 133(15):151801, 2024.

[168] Intae Yu. New Results from RENO. presented at Neutrino-2018. <https://zenodo.org/records/1287949>, June 2018.

[169] B. Jargowsky, L. Kolupaeva, and A. Yankelevich. Blessing Package For Ana2024 3F Bayesian Results Comparisons with Other Experiments. NOvA Internal Documents docid=62924, 2024.

[170] Patrick Dunne et al. Oscillation results from neutrino 2022, data release. Zenodo, <https://doi.org/10.5281/zenodo.6908532>, July 2022.

[171] Zoya Vallari. Results from a joint analysis of data from NOvA and T2K. Presented at FNAL seminar. <https://indico.fnal.gov/event/62062/>, Feb 2024.

[172] Edward Atkin. Results from the T2K+NOvA Joint Analysis. Presented at KEK seminar. <https://kds.kek.jp/event/49811/>, Feb 2024.

[173] K. Abe et al. (T2K). Measurements of neutrino oscillation parameters from the T2K experiment using  $3.6 \times 10^{21}$  protons on target. *Eur. Phys. J. C*, 83(9):782, 2023.